arXiv:1703.09531v1 [math.ST] 28 Mar 2017

A Bayesian nonparametric approach to log-concave density estimation

Ester Mariucci*, Kolyan Ray† and Botond Szabó†

Abstract

The estimation of a log-concave density on $\mathbb{R}$ is a canonical problem in the area of shape-constrained nonparametric inference. We present a Bayesian nonparametric approach to this problem based on an exponentiated Dirichlet process mixture prior and show that the posterior distribution converges to the log-concave truth at the (near-) minimax rate in Hellinger distance. Our proof proceeds by establishing a general contraction result based on the log-concave maximum likelihood estimator that prevents the need for further metric entropy calculations. We also present two computationally more feasible approximations and a more practical empirical Bayes approach, which are illustrated numerically via simulations.

AMS 2000 subject classifications: 62G07, 62G20.
Keywords: Density estimation, log-concavity, Dirichlet mixture, posterior distribution, convergence rate, nonparametric hypothesis testing.

1 Introduction

Nonparametric shape constraints offer practitioners considerable modelling flexibility by providing infinite-dimensional families that cover a wide range of parameters whilst also including numerous common parametric families. Log-concave densities on $\mathbb{R}$, that is densities whose logarithm is a concave function taking values in $[-\infty, \infty)$, constitute a particularly important shape-constrained class. This class includes many well-known parametric densities that are frequently used in statistical modelling, including the Gaussian, uniform, Laplace, Gumbel, logistic, gamma distributions with shape parameter at least one, Beta(α, β) distributions with $\alpha, \beta \geq 1$ and Weibull distributions with parameter at least one. Moreover, the class also preserves many of the attractive properties of Gaussian distributions, such as closure under

*Institut für Mathematik, Humboldt-Universität zu Berlin. This research was supported partly by the Federal Ministry for Education and Research through the Sponsorship provided by the Alexander von Humboldt Foundation and partly by the European Research Council under ERC Grant Agreement 320637. E-mail: mariucce@math.hu-berlin.de.

†Mathematical Institute, Leiden University. The research leading to these results has received funding from the European Research Council under ERC Grant Agreement 320637 and from the Netherlands Organization for Scientific Research (NWO) under Project number: 639.031.654.
E-mail: k.m.ray@math.leidenuniv.nl, b.t.szabo@math.leidenuniv.nl.

convolution, marginalization, conditioning and taking products. One can therefore view log-concave densities as a natural infinite-dimensional surrogate for Gaussians that retain many of their important features yet allow substantially more freedom, such as heavier tails. Log-concavity has been employed in economics [4, 5, 1], to investigate the convergence of Markov chain Monte Carlo sampling procedures [26, 3] and to propose new methodologies for many statistical problems, such as studying mixture models [38, 2], tail index estimation [27], clustering [7], regression [11] and independent component analysis [30]. For general reviews of inference with log-concave distributions and estimation under shape constraints, see [39] and [17] respectively.

The Bayesian approach provides a natural way to encode shape constraints via the prior distribution, for instance under monotonicity [32, 19, 29] or convexity contraints [33, 18]. We present here a Bayesian nonparametric method for log-concave density estimation on $\mathbb{R}$ based on an exponentiated Dirichlet process mixture prior, which we show converges to a log-concave truth in Hellinger distance at the (near-)minimax rate. To the best of our knowledge, this is the first Bayesian nonparametric approach to this problem. We also study two computationally motivated approximations to the full Dirichlet process mixture based on standard Dirichlet process approximations, namely the Dirichlet multinomial distribution and truncating the stick-breaking representation (see Chapter 4.3.3 of [15]). We further suggest an empirical Bayes approach that has clear practical advantages, while behaving similarly to the above in simulations. All of these priors are easily implementable using a random walk Metropolis-Hastings within Gibbs sampling algorithm, which we illustrate in Section 3.

An advantage of the Bayesian method is that point estimates and credible sets can be approximately computed as soon as one is able to sample from the posterior distribution. In particular, the posterior yields easy access to statements on Bayesian uncertainty quantification. The Bayesian approach also permits inference about multiple quantities, such as functionals of interest, in a unified way using the posterior distribution. For instance, there has been recent interest in estimating and constructing confidence intervals for the mode of a log-concave density [8] and the marginal posterior distribution of this quantity provides a natural approach to both questions. Whether this performs well, either theoretically or practically, is an interesting question that is, however, beyond the scope of this article. We also note that other constraints, such as a known mode [9], can similarly be enforced through suitable prior calibration.

Given the good performance of the log-concave MLE [10, 31, 6, 7, 11, 21, 20], one might expect that Bayesian procedures, being driven by the likelihood, behave similarly well. This is indeed the case, as we show below. Our proof relies on the classic testing approach of Ghosal et al. [13] with interesting modifications in the log-concave setting. The existence and optimality of the MLE in Hellinger distance is closely linked to a uniform control of bracketing entropy [35]. In our setting, one can exploit the affine equivariance of the log-concave MLE (Remark 2.4 of [11]) to circumvent the need to control the metric entropy of (almost) the whole space by reducing the problem to studying a subset satisfying restrictions on the first two moments of the underlying density. This is a substantial reduction, since obtaining sharp entropy bounds in even this reduced case is highly technical and requires significant effort, see Theorem 4 of Kim and Samworth [21]. One can then use the MLE to construct suitable plug-in tests with exponentially decaying type-II errors as in Giné and Nickl [16] that take

full advantage of the extra structure of the problem compared to the standard Le Cam-Birgé testing theory for the Hellinger distance [22]. Indeed, a naive attempt to control the entropy directly, as is standard in the Bayesian nonparametrics literature (e.g. [13]), results in an overly small set on which the prior must place most of its mass. This leads to unnecessary restrictions on the prior, which in particular are not satisfied by the priors we consider in Section 2, see Remark 1. Beyond this, there remain significant technical hurdles to proving that the prior places sufficient mass in a Kullback-Leibler neighbourhood of the truth, in particular related to the approximation of log-concave densities using piecewise log-linear functions with suitably spaced knots.

The paper is structured as follows. In Section 2 we introduce our priors and present our main results, both on general contraction for log-concave densities and for the specific priors considered here. In Section 3 we present a simulation study, including a more practical empirical Bayes implementation. In Section 4 we present the proofs of the main results with technical results placed in Section 5.

Notation: For two probability densities p and q with respect to Lebesgue measure λ on $\mathbb{R}$, we write $h^2(p,q) = \int(\sqrt{p}-\sqrt{q})^2$ for the squared Hellinger distance, $K(p,q) = \int p \log \frac{p}{q}$ for the Kullback-Leibler divergence and $V = \int p(\log \frac{p}{q})^2$. We denote by $P_{f_0}^n$ the product probability measure corresponding to the joint distribution of i.i.d random variables $X_1, \ldots, X_n$ with density f_0 and write $P_{f_0} = P_{f_0}^1$. For a function w, we denote by w'_- and w'_+ its left and right derivatives respectively, that is

$$w'_-(x) = \lim_{s \nearrow x} w'(s) \quad \text{and} \quad w'_+(x) = \lim_{s \searrow x} w'(s).$$

Let $\mathbb{R}^+ = [0, \infty)$ and for two real numbers a, b, let $a \wedge b$ and $a \vee b$ denote the minimum and maximum of a and b respectively. Finally, the symbols $\lesssim$ and $\gtrsim$ stand for an inequality up to a constant multiple, where the constant is universal or (at least) unimportant for our purposes.

2 Main Results

Consider i.i.d. density estimation, where we observe $X_1, ..., X_n \sim f_0$ with $f_0 = e^{w_0}$ an unknown log-concave density to be estimated. Let $\mathcal{F}$ denote the class of upper semi-continuous log-concave probability densities on $\mathbb{R}$. For $\alpha > 0$ and $\beta \in \mathbb{R}$, denote

$$\mathcal{F}_{\alpha,\beta} := \{f \in \mathcal{F} : f(x) \leq e^{\beta - \alpha|x|} \ \forall x \in \mathbb{R}\}.$$

By Lemma 1 of Cule and Samworth [6], for any log-concave density f_0 there exist constants $\alpha_{f_0} > 0$ and $\beta_{f_0} \in \mathbb{R}$ such that $f_0(x) \leq e^{\beta_{f_0} - \alpha_{f_0}|x|}$ for all $x \in \mathbb{R}$, and consequently any upper semi-continuous log-concave density is contained in $\mathcal{F}_{\alpha,\beta}$ for $0 < \alpha \leq \alpha_{f_0}$ and $\beta \geq \beta_{f_0}$. The sets $(\mathcal{F}_{\alpha,\beta})$ are thus natural sets for obtaining uniform statements in the context of log-concave density estimation.

We establish a general posterior contraction theorem for priors on log-concave densities using the general testing approach introduced in [13], which requires the construction of suitable tests with exponentially decaying type-II errors. We construct plug-in tests based

on the concentration properties of the log-concave MLE, similar to the linear estimators considered in [16, 28]. The MLE has been shown to converge to the truth at the minimax rate in Hellinger distance in Kim and Samworth [21] and the following theorem relies heavily on their result.

Theorem 1. *Let $\mathcal{F}$ denote the set of upper semi-continuous, log-concave probability densities on $\mathbb{R}$, let Π_n be a sequence of priors supported on $\mathcal{F}$. Consider a sequence $\varepsilon_n \to 0$ such that $n^{-2/5} \lesssim \varepsilon_n \lesssim n^{-\lambda}$ for some $3/8 < \lambda \leq 2/5$ and suppose there exists a constant $C > 0$ such that*

$$\Pi_n\Big(f \in \mathcal{F} : \int_{\mathbb{R}} f_0\Big(\log \frac{f_0}{f}\Big) \leq \varepsilon_n^2, \quad \int_{\mathbb{R}} f_0\Big(\log \frac{f_0}{f}\Big)^2 \leq \varepsilon_n^2\Big) \geq \exp(-Cn\varepsilon_n^2). \tag{1}$$

Then for sufficiently large M,

$$\Pi_n(f \in \mathcal{F} : h(f, f_0) \geq M\varepsilon_n | X_1, ..., X_n) \to 0$$

in $P_{f_0}^n$-probability.

The upper bound for ε_n is an artefact of the proof. Since our interest lies in obtaining the optimal rate $n^{-2/5}$, possibly up to logarithmic factors, it plays no further role in our results. It is typical in Bayesian nonparametrics to require metric entropy conditions, which come from piecing together tests for Hellinger balls into tests for the complements of balls, see for instance Theorem 7.1 of [13]. The lack of such a condition in Theorem 1 is tied to the optimality and specific structure of the log-concave MLE. Using the affine equivariance of the MLE (Remark 2.4 of [11]), one can reduce the testing problem to considering alternatives in the class $\mathcal{F}$ restricted to have zero mean and unit variance. Unlike the whole space $\mathcal{F}$, the bracketing Hellinger entropy of this latter set can be suitably controlled, thereby avoiding the need for additional entropy bounds.

Remark 1. *Obtaining sharp entropy bounds for log-concave function classes is a highly technical task and such bounds are only available for certain restricted subsets. Even in the case of mean and variance restrictions (Theorem 4 of [21]) and compactly supported and bounded densities (Proposition 14 of [20]), the proofs are lengthy and require substantial effort. To use such bounds for the classic entropy-based approach to prove posterior contraction would therefore require the prior to place most of its mass on the above types of restricted sets. For instance, the prior might be required to place all but exponentially small probability on $\mathcal{F}_{\alpha,\beta}$ for some given $\alpha > 0$, $\beta \in \mathbb{R}$. Such a prior construction is undesirable in practice and in fact none of our proposed priors satisfy such a restriction.*

We now introduce a prior on log-concave densities based on an exponentiated Dirichlet process mixture. For any measurable function $w : \mathbb{R} \to \mathbb{R}$, define the density

$$f_w(x) = \frac{e^{w(x)}}{\int_{\mathbb{R}} e^{w(y)} dy}, \tag{2}$$

which is well-defined if $\int_{\mathbb{R}} e^{w(y)} dy < \infty$. When w is a Gaussian process, such exponentiated priors have been considered both theoretically [36] and practically [24, 25]. Recall that any

monotone non-increasing probability density on $\mathbb{R}^+$ has a mixture representation [40]

$$f(x) = \int_x^\infty \frac{1}{u} dP(u),$$

where P is a probability measure on $\mathbb{R}^+$. Khazaei and Rousseau [19] and Salomond [29] used the above representation to obtain a Bayesian nonparametric prior for monotone non-increasing densities. Unfortunately, such a convenient mixture representation is unavailable for log-concave densities and so the prior construction is somewhat more involved. Integrating the above function and using that the (left or right) derivative of a concave function is monotone decreasing, consider the following concave function $w : \mathbb{R}^+ \to \mathbb{R}$,

$$w(x) = \gamma_1 \int_0^\infty \frac{u \wedge x}{u} dP(u) - \gamma_2 x,$$

where $\gamma_1 > 0$, $\gamma_2 \in \mathbb{R}$ and P is a probability measure on $[0, \infty)$. While not every concave function can be represented in this way, any log-concave density on $[0, \infty)$ can be approximated arbitrary well in Hellinger distance by a function of the form $e^w/(\int e^w)$, where w is as above with P a discrete probability measure, see Proposition 1. Translating the above thus gives a natural representation for a prior construction for log-concave densities on $\mathbb{R}$.

Consider therefore the following possibly n-dependent prior on the log-density $w : [a_n, b_n] \to \mathbb{R}$, where possibly $a_n \to -\infty$ and $b_n \to \infty$:

$$W(x) = \gamma_1 \int_0^{b_n - a_n} \frac{u \wedge (x - a_n)}{u} dP(u) - \gamma_2 (x - a_n), \tag{3}$$

where

- $P \sim DP(H \mathbb{1}_{[0, b_n - a_n]})$, the Dirichlet process with base measure $H \mathbb{1}_{[0, b_n - a_n]} = H(\mathbb{R}^+) \bar{H} \mathbb{1}_{[0, b_n - a_n]}$, where $0 < H(\mathbb{R}^+) < \infty$, $\bar{H}$ is a probability measure on $\mathbb{R}^+$ and every subset $U \subset [0, b_n - a_n]$ satisfies $H(U) \gtrsim \lambda(U)/(b_n - a_n)^\eta$ for some $\eta \geq 0$,
- $\gamma_i \sim p_{\gamma_i}$, $i = 1, 2$, where p_{γ_1}, p_{γ_2} are probability densities on $[0, \infty)$ and $\mathbb{R}$ respectively, satisfying $p_{\gamma_i}(|x|) \gtrsim e^{-c_i x^{1/4}}$, $c_i > 0$, for all $x \in [0, \infty)$ and $x \in \mathbb{R}$ respectively,
- γ_1, γ_2, and P are independent.

We denote by Π_n the full prior induced by f_W, where W is drawn as above. Some typical draws from the prior are plotted in Figure 1.

Remark 2. *If $(b_n - a_n)$ grows polynomially in n, then H must have polynomial tails. On the other hand, if $(b_n - a_n)$ grows more slowly than any polynomial, one can relax this condition. For instance, if H has a density h with respect to the Lebesgue measure, then it is sufficient that $\min_{t \in [0, b_n - a_n]} h(t) \gtrsim n^{-\lambda}$ for some $\lambda > 0$. In particular, if $(b_n - a_n) \lesssim \log n$, then h may have exponential tails.*

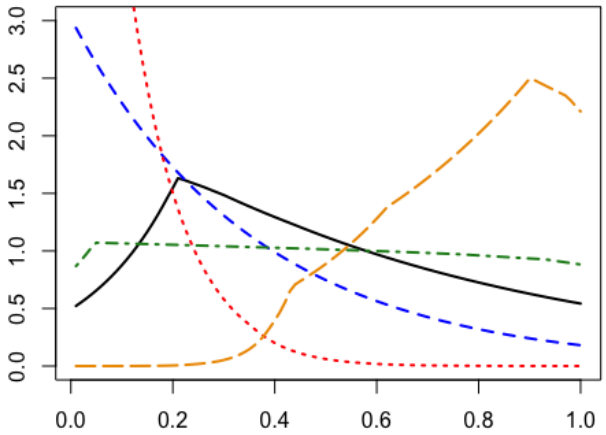

Figure 1: Prior draws with $[a_n, b_n] = [0, 1]$, $\gamma_1 \sim \text{Cauchy}_+(0, 1)$, $\gamma_2 \sim \text{Cauchy}(0, 1)$, $H = U(0, 1)$ and using the stick breaking construction.

Let us comment on several aspects of our prior. Firstly, since Dirichlet process draws are atomic with probability one, the prior draws (3) will be piecewise linear and concave. Moreover, we could add any concave function to (3), such as an $-\gamma_3 x^2$-type term, and still have a suitable concave prior. This permits greater modelling flexibility but complicates computation. In any case, the prior described above gives optimal contraction rates and can be computed in practice, so we restrict our attention to it. Another point to note is that if $(b_n - a_n) \to \infty$ and H is supported on the whole of $\mathbb{R}^+$, then the Dirichlet process base measure has total mass $H(\mathbb{R}^+)\bar{H}([a_n, b_n]) \leq H(\mathbb{R}^+)$ for fixed n. This has the interpretation of assigning the prior more weight as $n \to \infty$, up to the full prior weight $H(\mathbb{R}^+)$. An alternative would be to re-weight the base measure to have full mass $H(\mathbb{R}^+)$ to give it equal weight for all n. This plays no role asymptotically and so we restrict to the first case for technical convenience.

A potentially more serious issue is that for fixed n, the support of the prior draws may not contain the support of the true density f_0, in which case observations outside $[a_n, b_n]$ cause the likelihood to be identically zero. While this is not a problem for n large enough if $-a_n, b_n \to \infty$ fast enough (see Theorem 3), it can be an issue for finite n. In practice, if one has an idea of the support of f_0, it is enough to select $[a_n, b_n]$ large enough to contain $\text{supp}(f_0)$. A more pragmatic solution is to use an empirical Bayes approach and make the prior data-dependent by setting $a_n := X_{(1)}$, $b_n := X_{(n)}$ the first and last order statistics. This ensures that the likelihood is never zero and the posterior is always well-defined. Indeed, the MLE is supported on $[X_{(1)}, X_{(n)}]$ and so this can be thought of as plugging-in an estimate of

the approximate support based on the likelihood. Moreover, since this approach yields the smallest support $[a_n, b_n]$ with non-zero likelihood, it also brings computational advantages. In particular, it can prevent the need to simulate the posterior distribution on potentially very large regions of $\mathbb{R}$ where the posterior draws are essentially indistinguishable from zero. The empirical Bayes method behaves very similarly to the prior (3) in simulations and is the approach we would advocate in practice. Further practical considerations are discussed below.

We first present a contraction result when the true density f_0 has known compact support.

Theorem 2. *Let $f_0 \in \mathcal{F}_{\alpha,\beta}$ for some $\alpha > 0$, $\beta \in \mathbb{R}$ and suppose further that f_0 is compactly supported. Let $a_n \equiv a$ and $b_n \equiv b$ for all n and denote by $\Pi_n = \Pi$ the prior described above. If $\mathrm{supp}(f_0) \subset [a, b]$, then*

$$\Pi(f : h(f, f_0) \geq M(\log n)n^{-2/5} \mid X_1, ..., X_n) \to 0$$

in $P_{f_0}^n$-probability for some $M = M(\alpha, \beta) > 0$.

If $\mathrm{supp}(f_0)$ is not contained in a compact set or is unknown, it suffices to let $-a_n, b_n \to \infty$ fast enough. A slightly stronger lower bound on the tail of p_{γ_1} is consequently required, depending on the size of $(b_n - a_n)$.

Theorem 3. *Let $f_0 \in \mathcal{F}_{\alpha,\beta}$ for some $\alpha > 0$, $\beta \in \mathbb{R}$ and let Π_n denote the prior described above with $-a_n, b_n \gg \log n$. Assume further that $(b_n - a_n) \lesssim n^{\mu/5}$ and that the prior density p_{γ_1} for γ_1 satisfies the stronger lower bound $p_{\gamma_1}(x) \gtrsim e^{-c_1 x^{1/(4+\mu)}}$ for some $0 \leq \mu < 4$. Then*

$$\Pi_n(f : h(f, f_0) \geq M\varepsilon_n \mid X_1, ..., X_n) \to 0$$

in $P_{f_0}^n$-probability for some $M = M(\alpha, \beta) > 0$ and

$$\varepsilon_n = \max\left((\log n)n^{-2/5}, (b_n - a_n)n^{-4/5}\right). \tag{4}$$

Theorem 2 follows immediately from Theorem 3 and so its proof is omitted. If $(b_n - a_n) = O((\log n)n^{2/5})$, then we obtain the minimax rate for log-concave density estimation in Theorem 3, up to a logarithmic factor. Note also that the above implies posterior convergence in total variation at the same rate ε_n given in (4). We also mention that all the above statements are proved uniformly over $f_0 \in \mathcal{F}_{\alpha,\beta}$.

Dirichlet process mixture priors are popular in density estimation due to the conjugacy of the posterior distribution, thereby providing methods that are highly efficient computationally. However, due to the exponentiation (2), this conjugacy property no longer holds, resulting in a less attractive prior choice that brings computational challenges. In practice, it is common to use approximations of the Dirichlet process to speed up computations, see for instance Chapter 4.3.3 of [15].

We firstly consider the Dirichlet multinomial distribution as a replacement for the Dirichlet process in our prior. By the proof of Theorem 3, the underlying true log-concave density can be well approximated by a piecewise log-linear density with at most $N = Cn^{1/5} \log n$ knots,

for some large enough constant $C > 0$. In view of this, it is reasonable to take N atoms in the distribution. The corresponding prior on log-concave densities then takes the form

$$\begin{aligned} \theta_i &\overset{iid}{\sim} \bar{H}\mathbb{1}_{[0,b_n-a_n]}, \quad \text{for } i = 1, ..., N, \\ p = (p_1, ...p_N) &\sim Dir(\alpha_1, ..., \alpha_N), \\ \gamma_i &\overset{iid}{\sim} p_{\gamma_i}, \quad i = 1, 2, \\ f_{\theta,p,\gamma_1,\gamma_2}(x) &= \frac{\exp\{\gamma_1 \sum_{i=1}^N \frac{\theta_i \wedge (x-a_n)}{\theta_i} p_i - \gamma_2(x-a_n)\}\mathbb{1}_{[a_n,b_n]}(x)}{\int_{a_n}^{b_n} \exp\{\gamma_1 \sum_{i=1}^N \frac{\theta_i \wedge (u-a_n)}{\theta_i} p_i - \gamma_2(u-a_n)\}du}, \end{aligned} \tag{5}$$

where α_i, $i = 1, ..., N$, are chosen such that $\alpha_i \geq N^{-\nu}$ for some arbitrary large $\nu > 0$ and $\sum_i \alpha_i \leq H(\mathbb{R}^+)$.

An alternative choice for the mixing prior is to truncate the stick-breaking representation of the Dirichlet process at a fixed level. Similarly to the Dirichlet multinomial distribution, we truncate the stick-breaking process at level $N = Cn^{1/5} \log n$, resulting in the same hierarchical prior as in (5) with the only difference being that the distribution of p in the N-simplex is given by

$$p_i \sim V_i \prod_{j=1}^{i-1}(1 - V_j), \text{ where } V_i \sim \text{Beta}(1, H(\mathbb{R}^+)), \; i = 1, ..., N-1. \tag{6}$$

Both of these computationally more efficient approximations have the same theoretical guarantees as the full exponentiated Dirichlet process prior Π_n.

Corollary 1. *Let $f_0 \in \mathcal{F}_{\alpha,\beta}$ for some $\alpha > 0$, $\beta \in \mathbb{R}$ and let Π_n' denote either the prior (5) or (6). If $-a_n, b_n \gg \log n$, $(b_n - a_n) \lesssim n^{\mu/5}$ and the prior density p_{γ_1} for γ_1 satisfies the stronger lower bound $p_{\gamma_1}(x) \gtrsim e^{-c_1 x^{1/(4+\mu)}}$ for some $0 \leq \mu < 4$, then*

$$\Pi_n'(f : h(f, f_0) \geq M\varepsilon_n \mid X_1, ..., X_n) \to 0$$

in $P_{f_0}^n$-probability for some $M = M(\alpha, \beta) > 0$ and ε_n given by (4).

The proofs of Theorem 3 and Corollary 1 establish the small-ball probability (1) by approximating a log-concave density in $\mathcal{F}_{\alpha,\beta}$ with a suitable piecewise log-linear density. This approximation requires several key properties, which makes its construction non-standard and technically involved, and it may be of independent interest. The proof of Proposition 1 is deferred to Section 5.1.

Proposition 1. *Let $f_0 \in \mathcal{F}_{\alpha,\beta}$ and $([a_n, b_n])_n$ be a sequence of compact intervals such that $[-\frac{8}{5\alpha} \log n, \frac{8}{5\alpha} \log n] \subset [a_n, b_n]$ and $(b_n - a_n) = o(n^{4/5})$. For any $n \geq n_0(\alpha, \beta)$, there exists a log-concave density $\bar{f}_n$ that is piecewise log-linear with $\bar{N} = O(n^{1/5} \log n)$ knots $z_1, ..., z_{\bar{N}} \in [0, b_n - a_n]$ and satisfies the following properties:*

(i) $h^2(f_0, \bar{f}_n) \leq C[(\log n)^2 n^{-4/5} + (b_n - a_n)^2 n^{-8/5}]$,

(ii) $\{x \in \mathbb{R} : \bar{f}_n(x) > 0\} = [a_n, b_n]$,

(iii) *the knots are* $cn^{-6/5}\log n$*-separated for some* $c > 0$,

(iv) $f_0(x) \leq C\bar{f}_n(x)$ *for all* $x \in [a_n, b_n]$,

(v) *there exist* $\bar{\gamma}_1 \in [0, 2(b_n - a_n)n^{4/5}]$, $|\bar{\gamma}_2| \leq n^{4/5}$, $\bar{\gamma}_3 \in \mathbb{R}$ *and* $(\bar{p}_1, .., \bar{p}_{\bar{N}})$ *satisfying* $p_i \geq 0$ *and* $\sum_{i=1}^{\bar{N}} p_i = 1$, *such that*

$$\bar{f}_n(x) = \exp\left(\bar{\gamma}_1 \sum_{i=1}^{\bar{N}} \frac{z_i \wedge (x - a_n)}{z_i}\bar{p}_i - \bar{\gamma}_2(x - a_n) + \bar{\gamma}_3\right) \mathbb{1}_{[a_n, b_n]}(x).$$

It is relatively straightforward to establish an approximation of f_0 satisfying (i). However, approximating f_0 by $\bar{f}_n$ in a Kullback-Leibler type sense, as in (1), necessitates control of the support of $\bar{f}_n$ via (ii) and uniform control of the ratio $f_0/\bar{f}_n$ via (iv). The most difficult property to establish is the polynomial separation of the points in (iii). This is needed to ensure that the Dirichlet process prior simultaneously puts sufficient mass in a neighbourhood of each of the knots z_i, $i = 1, \ldots, \bar{N}$. Setting $[a_n, b_n] = [-\frac{8}{5\alpha}\log n, \frac{8}{5\alpha}\log n]$ yields the following corollary.

Corollary 2. *Let* $f_0 \in \mathcal{F}_{\alpha,\beta}$. *For any* $n \geq n_0(\alpha, \beta)$, *there exists a log-concave density* $\bar{f}_n$ *supported on* $[-\frac{8}{5\alpha}\log n, \frac{8}{5\alpha}\log n]$ *that is piecewise log-linear with* $O(n^{1/5}\log n)$ *knots and satisfies* $h^2(f_0, \bar{f}_n) \leq C(\log n)^2 n^{-4/5}$. *Moreover, we may take the knots to be* $cn^{-6/5}\log n$*-separated for some* $c > 0$.

3 Simulation study

We present a simulation study to assess the performance of the proposed log-concave priors for density estimation. In particular, we investigate the prior based on the truncated stick breaking representation (6), firstly with deterministically chosen support $[a_n, b_n]$ and secondly its empirical Bayes counterpart with support $[X_{(1)}, X_{(n)}]$, where $X_{(1)}$ and $X_{(n)}$ denote the smallest and largest observations, respectively.

Consider first the posterior distribution arising from the prior with deterministic support $[a_n, b_n]$. We have drawn random samples of size $n = 50, 200, 500$ and 2500 from a gamma distribution with shape and rate parameters two and one, respectively. We took the number of linear pieces in the exponent of the prior to be $m = n^{1/5}\log n$, set $[a_n, b_n] = [-(\log n)^{1+\epsilon}, (\log n)^{1+\epsilon}]$ for $\epsilon > 0$ small, endowed the break-point parameters $\theta = (\theta_1, ..., \theta_m)$ with independent uniform priors on $[0, b_n - a_n]$, assigned the weight parameters p a stick-breaking distribution truncated at level m, and endowed γ_1 and γ_2 a half Cauchy and a Cauchy distribution, respectively, with location parameter zero and scale parameter 1. Since the posterior distribution does not have a closed-form expression, we drew approximate samples from the posterior using a random walk Metropolis-Hastings within Gibbs sampling algorithm with a burn-in period of 3000 iterations, before running the algorithm for a further 3000 iterations. In Figure 2, we have plotted the true distribution with a solid red line, the posterior mean with a solid blue line and the 95% pointwise-credible band with dashed blue lines. The data is represented by a histogram on the figures. We see that the posterior mean

gives an adequate estimator for the true log-concave density and the 95% credible bands mostly contain the true function, except for points close to zero. It should however be noted that the frequentist coverage of Bayesian credible sets is a delicate subject in high-dimensions, see for instance Szabó et al. [34], and is beyond the scope of this article.

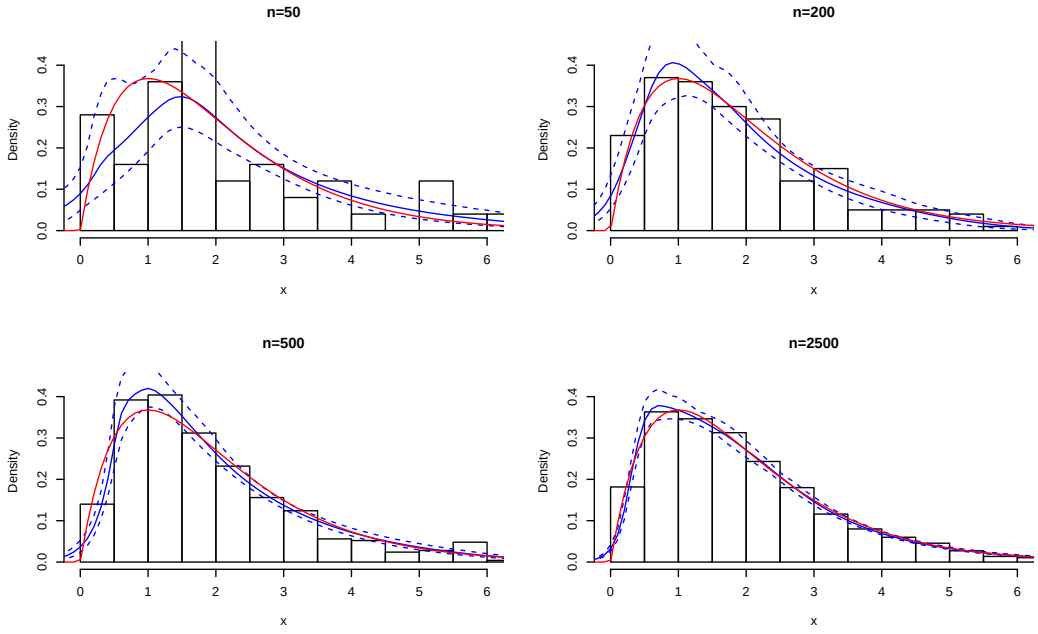

Figure 2: Prior with $[a_n, b_n]$ selected deterministically: the underlying Gamma(2,1) density function (red), posterior mean (solid blue), point-wise credible bands (dashed blue) and data is represented with a histogram. We have increasing sample size from left to right and top to bottom $n = 50, 200, 500$ and 2500.

We next investigate the behaviour of the empirical Bayes version of the proposed prior. We again simulated $n = 50, 200, 500$ and 2500 independent observations from a Gamma(2,1) distribution and set the compact support of the prior densities to be $[a_n, b_n] = [X_{(1)}, X_{(n)}]$, the smallest and largest observations. As in the non-data dependent prior, we set $m = n^{1/5} \log n$ and endowed the parameters θ, p, γ_1 and γ_2 with the same priors as above. We ran the algorithm for the same number of iterations, again taking the first half of the chain as a burn-in period. We plotted the outcome in Figure 3. One can see that for $n \geq 500$ observations, the posterior mean (solid blue) closely resembles the underlying true gamma density (solid red), while the fit is already reasonable for $n = 200$. The pointwise 95%-credible bands contain the true density, even near zero, which was problematic in case of the prior with support selected deterministically. Comparing Figures 2 and 3, we see that the empirical Bayes approach of selecting the support $[a_n, b_n]$ in a data-driven way outperforms a deterministic selection. We also note that the algorithm for the empirical Bayes method was considerably faster due to the smaller support, which reduces the computation time of the normalizing constants $\int e^{w(y)} dy$

of the densities.

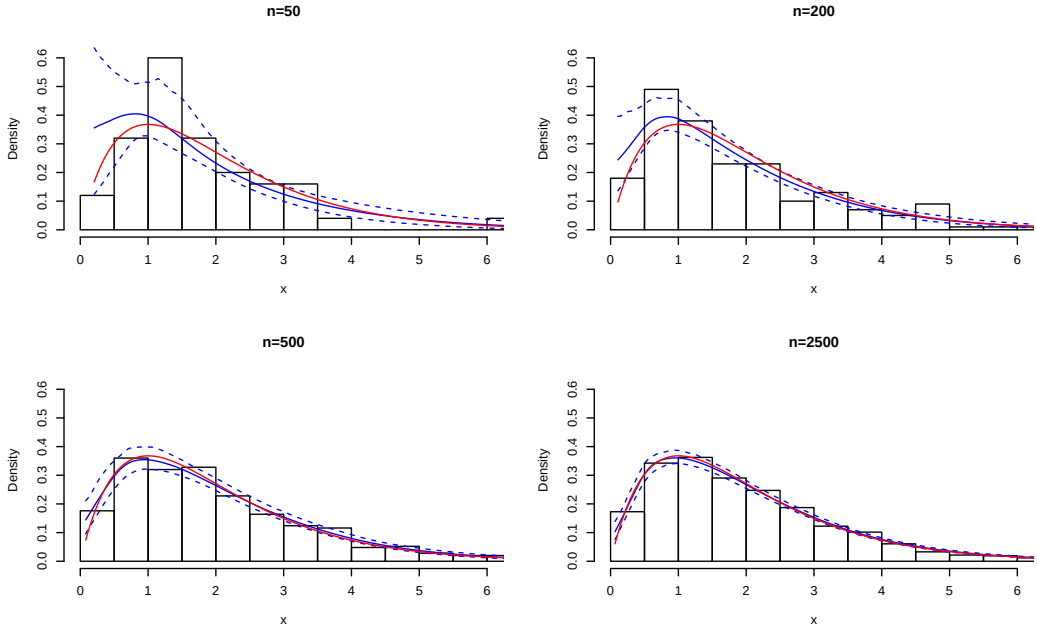

Figure 3: Prior with data-driven support: the underlying Gamma(2,1) density function (red), posterior mean (solid blue), point-wise credible bands (dashed blue) and data is represented with a histogram. We have increasing sample size from left to right and top to bottom $n = 50, 200, 500$ and 2500.

We lastly investigate the performance of the posterior distribution corresponding to the data-driven version of the prior for recovering different underlying density functions. We have considered a standard normal distribution, a gamma distribution with shape parameter 2 and rate parameter 1, a beta distribution with shape parameters 2 and 3, and a Laplace distribution with location parameter 0 and dispersion parameter 1. In all four examples we have taken sample size $n = 1500$. The posterior mean (solid blue) and 95% point-wise credible bands (dashed blue) are visualized in Figure 4. All four pictures show satisfactory results, both for estimation using the posterior mean and for coverage using the point-wise credible band. We note that the displayed plots convey typical behaviour and are representative of multiple simulations. We hence draw the conclusion that the proposed method seems to work well in practice for various choices of common log-concave densities.

4 Proofs

Define the following classes of log-concave densities with mean and variance restrictions:

$$\overline{\mathcal{F}}^{\xi,\eta} = \left\{ f \in \mathcal{F} : \mu_f := \int x f(x) dx = \xi, \quad \sigma_f^2 := \int (x - \mu_f)^2 f(x) dx = \eta \right\}$$

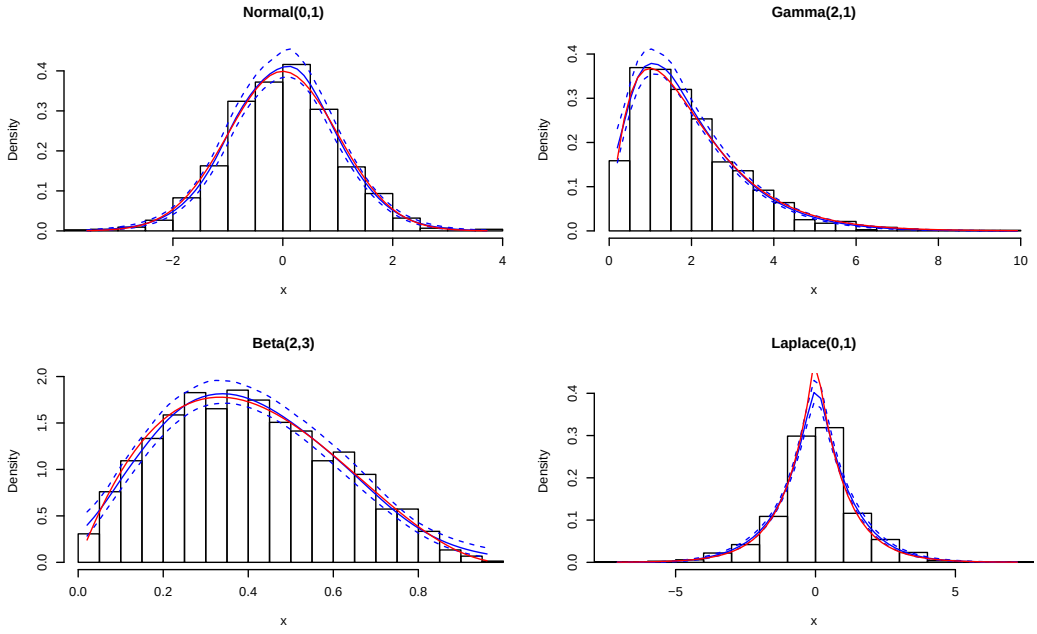

Figure 4: The underlying density function (red), the posterior mean (solid blue), the pointwise credible bands (dashed blue). The data is represented with a histogram. The true density functions are from left to right and top to bottom: standard Gaussian, Gamma(2,1), Beta(2,3) and Laplace(0,1).

and

$$\widetilde{\mathcal{F}}^{\xi,\eta} = \{f \in \mathcal{F} : |\mu_f| \leq \xi, \quad |\sigma_f^2 - 1| \leq \eta\}.$$

Let $\hat{f}_n$ denote the log-concave MLE based on i.i.d. random variables $X_1, \ldots, X_n$ arising from a density $f_0 \in \mathcal{F}$.

The proof of Theorem 1 relies on a concentration inequality for the log-concave MLE based on data from moment-restricted densities. This is the content of the following lemma, whose proof is essentially contained in Kim and Samworth [21] for the more difficult case of general $d \geq 1$. However, we require a sharper probability bound than they provide and so make some minor modifications to their argument. The proof is deferred to Section 5.2

Lemma 1. *For every $\varepsilon > 0$, there exist positive constants L_0, $C(\varepsilon)$ and $c(\varepsilon)$ such that for any $L \geq L_0$ and $n \geq n_0(L)$,*

$$\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n \left(h(\hat{g}_n, g_0) \geq L n^{-2/5} \right) \leq C(\varepsilon) \exp\left(-c(\varepsilon) L^2 n^{1/(4+2\varepsilon)} \right),$$

where $\hat{g}_n$ denotes the log-concave maximum likelihood estimator based on an i.i.d. sample $Z_1, \ldots, Z_n$ from g_0.

Proof of Theorem 1. As in the proof of Theorem 2.1 of [13], using the lower bound on the small ball probability from (1), it is sufficient to construct tests $\phi_n = \phi_n(X_1, ..., X_n; f_0)$ such that

$$P_{f_0}^n \phi_n \to 0, \qquad \text{and} \qquad \sup_{f \in \mathcal{F}: h(f,f_0) \geq M\varepsilon_n} P_f^n(1-\phi_n) \leq e^{-(C+4)n\varepsilon_n^2}$$

for n large enough, where the constant $C > 0$ matches that in (1).

For M_0 a constant to be chosen below, set $\phi_n = \mathbb{1}\{h(\hat{f}_n, f_0) \geq M_0\varepsilon_n\}$, where $\hat{f}_n$ is the log-concave MLE based on i.i.d. observations $X_1, ..., X_n$ from a density $f_0 \in \mathcal{F}$. Let $\mu_{f_0} = \mathbb{E}X_i$, $\sigma_{f_0}^2 = \mathrm{Var}(X_i)$ and define $Z_i = (X_i - \mu_{f_0})/\sigma_{f_0}$, so that $\mathbb{E}Z_i = 0$ and $\mathrm{Var}(Z_i) = 1$. Further set $g_0(z) = \sigma_{f_0} f_0(\sigma_{f_0} z + \mu_{f_0})$ and $\hat{g}_n(z) = \sigma_{f_0}\hat{f}_n(\sigma_{f_0} z + \mu_{f_0})$, so $g_0 \in \overline{\mathcal{F}}^{0,1}$. By affine equivariance (Remark 2.4 of [11]), $\hat{g}_n$ is the log-concave maximum likelihood estimator of g_0 based on $Z_1, ..., Z_n$.

Using Lemma 1 and the invariance of the Hellinger distance under affine transformations, the type-I error satisfies

$$P_{f_0}^n \phi_n = P_{g_0}^n(h(\hat{g}_n, g_0) \geq M_0\varepsilon_n) \leq P_{g_0}^n(h(\hat{g}_n, g_0) \geq L_0 n^{-2/5}) \to 0$$

as $n \to \infty$ for M_0 large enough. For $f \in \mathcal{F}$ such that $h(f, f_0) \geq M\varepsilon_n$,

$$\begin{aligned} P_f^n(1-\phi_n) &= P_f^n(h(f_0, \hat{f}_n) < M_0\varepsilon_n) \\ &\leq P_f^n(h(f_0, f) - h(f, \hat{f}_n) < M_0\varepsilon_n) \\ &\leq P_f^n((M - M_0)\varepsilon_n < h(f, \hat{f}_n)). \end{aligned}$$

Letting $L > 0$ be an arbitrarily large constant and $\epsilon(\lambda) = \frac{8\lambda - 3}{2 - 4\lambda} > 0$, we can take $M > 0$ large enough so that applying Lemma 1,

$$\begin{aligned} \sup_{f \in \mathcal{F}:\, h(f,f_0) \geq M\varepsilon_n} P_f^n(1-\phi_n) &\leq \sup_{g \in \overline{\mathcal{F}}^{0,1}} P_g^n(h(\hat{g}_n, g) \geq (M - M_0)\varepsilon_n) \\ &\leq \sup_{g \in \overline{\mathcal{F}}^{0,1}} P_g^n(h(\hat{g}_n, g) \geq (M - M_0) n^{-2/5}) \\ &\leq C(\epsilon) e^{-c(\epsilon)(M - M_0)^2 n^{1/(4+2\epsilon)}} \leq e^{-Ln\varepsilon_n^2} \end{aligned}$$

for $n \geq n_0(L)$ large enough. □

Proof of Theorem 3. Let $f_0 \in \mathcal{F}_{\alpha,\beta}$ for some $\alpha > 0$ and $\beta \in \mathbb{R}$. We may restrict to a suitable compactly supported density approximating f_0 using the first paragraph of the proof of Theorem 2 of Ghosal and van der Vaart [14]. For completeness we reproduce their argument in this paragraph. Let $\phi(x) = \mathbb{1}_{[-1,1]}(x)$ and set $\phi_n(x) = \phi(x/t_n)$ for $t_n = a' \log n$ for some $a' > \alpha^{-1}$. Define new observations $\bar{X}_1, ..., \bar{X}_{\bar{n}}$ from the original observations $X_1, ..., X_n$ by rejecting each X_i independently with probability $1 - \phi_n(X_i)$. Since $P_{f_0}[-t_n, t_n]^c \leq 2e^{\beta}\alpha^{-1}e^{-\alpha t_n} = o(n^{-1})$, the probability that at least one of the X_i's is rejected is $o(1)$ and so the posterior based on the original and modified observations are the same with $P_{f_0}^n$-probability tending to one. Since posterior contraction is defined via convergence in $P_{f_0}^n$-probability, this implies that the posterior contraction rates are the same. The new observations come from a density

$f_{0,n}$ that is proportional to $f_0\phi_n$, which is log-concave and upper semi-continuous. Since $|1-\int f_0\phi_n| \le P_{f_0}[-t_n,t_n]^c = o(n^{-1})$,

$$\begin{aligned} h^2(f_0,f_{0,n}) &\le 2\int_{\mathbb{R}} f_0\left(1-\frac{1}{\sqrt{\int f_0\phi_n}}\right)^2 dx + \frac{2}{\int f_0\phi_n}\int_{\mathbb{R}} f_0(1-\sqrt{\phi_n})^2 dx \\ &\le 2\frac{(\int f_0\phi_n - 1)^2}{\int f_0\phi_n} + \frac{2}{\int f_0\phi_n}\int_{\mathbb{R}\setminus[-t_n,t_n]} f_0 dx = o(n^{-1}). \end{aligned} \tag{7}$$

It is therefore sufficient to establish contraction for the posterior based on the new observations about the density $f_{0,n} = f_0\phi_n/\int f_0\phi_n$.

We apply Theorem 1 so that we only need to show the small-ball probability (1). Note that $f_{0,n}(x) \le e^{\beta-\alpha|x|}(1+o(n^{-1}))$, so that $f_{0,n} \in \mathcal{F}_{\alpha,2\beta}$ for n large enough. Since $-a_n, b_n \gg \log n$ and $b_n - a_n = o(n^{4/5})$, we may construct an approximation $\bar{f}_n$ of $f_{0,n}$ based on the interval $[a_n,b_n]$ for n large enough using Proposition 1. By Lemma 8 of [14],

$$\begin{aligned} \int_{\mathbb{R}} f_{0,n}\Big(\log\frac{f_{0,n}}{f_W}\Big)^k &\lesssim \big(h^2(f_{0,n},\bar{f}_n) + h^2(\bar{f}_n,f_W)\big) \\ &\quad\times\left(1+\log\left\|\frac{f_{0,n}}{\bar{f}_n}\right\|_{L^\infty([a_n,b_n])} + \log\left\|\frac{\bar{f}_n}{f_W}\right\|_{L^\infty([a_n,b_n])}\right)^k \end{aligned}$$

for $k=1,2$. By Proposition 1(i) and (iv), the first term in the first bracket and the second term in the second bracket are $O((\log n)^2 n^{-4/5} + (b_n-a_n)^2 n^{-8/5})$ and $O(1)$ respectively.

By Proposition 1(v), $\bar{f}_n$ has representation

$$\bar{f}_n(x) = \exp\left(\bar{\gamma}_1\sum_{i=1}^{\bar{N}}\frac{z_i\wedge(x-a_n)}{z_i}\bar{p}_i - \bar{\gamma}_2(x-a_n) + \bar{\gamma}_3\right)\mathbb{1}_{[a_n,b_n]}(x), \tag{8}$$

where $(z_i)_{i=1}^{\bar{N}} \subset [0,b_n-a_n]$ are the knots written in increasing order, $\bar{N} = \bar{N}_n = O(n^{1/5}\log n)$ and $\sum_{i=1}^{\bar{N}}\bar{p}_i = 1$. Let $\bar{w}_n(x) = (\log\bar{f}_n(x) - \bar{\gamma}_3)\mathbb{1}_{[a_n,b_n]}(x) - \infty\mathbb{1}_{\mathbb{R}\setminus[a_n,b_n]}(x)$ so that $\bar{f}_n = f_{\bar{w}_n}$ using the transformation (2). We may thus without loss of generality take $\bar{\gamma}_3 = 0$ since it is contained in the normalization (2).

Suppose that f_w is a (log-concave) density with support equal to $[a_n,b_n]$ and such that $\|\bar{w}_n - w\|_{L^\infty([a_n,b_n])} \le c\varepsilon_n$. Since $\int e^w = e^{O(\varepsilon_n)}\int e^{\bar{w}_n}$, it follows that for $x\in[a_n,b_n]$, $\bar{f}_n(x)/f_w(x) \le e^{O(\varepsilon_n)} = e^{o(1)}$. Since $h^2(\bar{f}_n,f_w) \lesssim \varepsilon_n^2$ by Lemma 3.1 of [36], we can conclude that

$$\{w : \|\bar{w}_n - w\|_{L^\infty([a_n,b_n])} \le c\varepsilon_n\} \subset \{w : K(f_{0,n},f_w) \le \varepsilon_n^2,\ V(f_{0,n},f_w) \le \varepsilon_n^2\} \tag{9}$$

for some $c>0$. It therefore suffices to lower bound the prior probability of the left-hand set.

Fix $\delta > 0$ to be chosen sufficiently large below. Since the z_i are $n^{-6/5}$-separated by Proposition 1(iii), we can find a collection of disjoint intervals $(U_i)_{i=1}^{\bar{N}}$ in $[a_n,b_n]$ such that $\lambda(U_i) = \varepsilon_n^\delta$ and $z_i \in U_i$ for $i=1,...,\bar{N}$. Further denote $U_0 := \mathbb{R}\setminus\cup_{i=1}^{M_n} U_i$. Let W be a prior

draw of the form (3) with parameters γ_1, γ_2 and P. Writing $p_i = P(U_i)$, $\bar{p}_0 = 0$ and using the triangle inequality, for any $x \in [a_n, b_n]$,

$$\begin{aligned}
|\bar{w}_n(x) - W(x)| &= \left| \bar{\gamma}_1 \sum_{i=1}^{\bar{N}} \frac{z_i \wedge (x - a_n)}{z_i} \bar{p}_i - \bar{\gamma}_2 (x - a_n) - \gamma_1 \int_0^\infty \frac{\theta \wedge (x - a_n)}{\theta} dP(\theta) - \gamma_2 (x - a_n) \right| \\
&\leq |\bar{\gamma}_1 - \gamma_1| \int_0^\infty \frac{\theta \wedge (x - a_n)}{\theta} dP(\theta) + \bar{\gamma}_1 \left| \int_{U_0} \frac{\theta \wedge (x - a_n)}{\theta} dP(\theta) \right| \\
&\quad + \bar{\gamma}_1 \left| \sum_{i=1}^{\bar{N}} \int_{U_i} \frac{\theta \wedge (x - a_n)}{\theta} dP(\theta) - \sum_{i=1}^{\bar{N}} \frac{z_i \wedge (x - a_n)}{z_i} p_i \right| \\
&\quad + \bar{\gamma}_1 \left| \sum_{i=1}^{\bar{N}} \frac{z_i \wedge (x - a_n)}{z_i} (p_i - \bar{p}_i) \right| + (b_n - a_n) |\bar{\gamma}_2 - \gamma_2| \\
&\leq |\bar{\gamma}_1 - \gamma_1| + \bar{\gamma}_1 \sum_{i=1}^{\bar{N}} \sup_{\theta \in U_i} \frac{|\theta - z_i|}{\theta \wedge z_i} p_i + \bar{\gamma}_1 p_0 + \bar{\gamma}_1 \sum_{i=1}^{\bar{N}} |p_i - \bar{p}_i| + (b_n - a_n) |\bar{\gamma}_2 - \gamma_2| \\
&\leq |\bar{\gamma}_1 - \gamma_1| + 2\bar{\gamma}_1 \sum_{i=1}^{\bar{N}} \frac{\lambda(U_i)}{z_i} p_i + \bar{\gamma}_1 \sum_{i=0}^{\bar{N}} |p_i - \bar{p}_i| + (b_n - a_n) |\bar{\gamma}_2 - \gamma_2|, \qquad (10)
\end{aligned}$$

where we have used in the second to last line that the maximal distance between the (piecewise) lines $(y \wedge a)/a$ and $(y \wedge b)/b$ occurs at $y = a \wedge b$ and in the last line that $\theta > z_i/2$ for all $\theta \in U_i$ for sufficiently large choice of the parameter $\delta > 0$. By Proposition 1(v), we have $\bar{\gamma}_1 \lesssim n^{4/5}(b_n - a_n) \lesssim n^{(4+\mu)/5}$. Furthermore, by the separation of the knots, $z_i \geq z_1 \geq cn^{-6/5}$, $i = 1, ..., \bar{N}$, and so by the assumptions on the (U_i), the second term is bounded by $2c^{-1}\bar{\gamma}_1 n^{6/5} \varepsilon_n^\delta \leq \tilde{c}\varepsilon_n$ for some $\delta, \tilde{c} > 0$ large enough.

The remaining three terms are independent under the prior and so can be dealt with separately. By the assumptions on the base measure of the Dirichlet process, we have that $\sum_{i=0}^{\bar{N}} H(U_i) \leq H(\mathbb{R}^+)$ and $H(U_i) = H(\mathbb{R}^+)\bar{H}(U_i) \gtrsim \lambda(U_i)/(b_n - a_n)^\eta \geq \varepsilon_n^\delta/(b_n - a_n)^\eta \geq \varepsilon_n^{\delta'}$ for $i = 1, ..., \bar{N}$ and some $\delta' > \delta$. For $i = 0$, note that $\lambda(U_0) \geq (b_n - a_n) - \bar{N}\varepsilon_n^\delta \gtrsim 1$. Using the lower the bounds for the $\lambda(U_i)$, which come from the polynomial separation of the knots in Proposition 1(iii), we can apply Lemma 10 of [14] to get

$$\Pi_n \left(\bar{\gamma}_1 \sum_{i=0}^{\bar{N}} |p_i - \bar{p}_i| \leq \varepsilon_n \right) \gtrsim e^{-c\bar{N} \log(2\bar{\gamma}_1/\varepsilon_n)} \gtrsim e^{-c' n \varepsilon_n^2}. \qquad (11)$$

From the tail assumption on the density of γ_1 and the upper bound on $\bar{\gamma}_1$, we have

$$\Pi_n\big(|\gamma_1 - \bar{\gamma}_1| \leq \varepsilon_n\big) \gtrsim \varepsilon_n e^{-c(\bar{\gamma}_1 + \varepsilon_n)^{1/(4+\mu)}} \geq e^{-c' n^{1/5}} \geq e^{-n\varepsilon_n^2}. \qquad (12)$$

By Proposition 1(v), $|\bar{\gamma}_2| \lesssim n^{4/5}$, which, combined with the tail bound on the density of γ_2, yields

$$\Pi_n((b_n - a_n)|\gamma_2 - \bar{\gamma}_2| \leq \varepsilon_n) \gtrsim \tfrac{\varepsilon_n}{b_n - a_n} e^{-c(|\bar{\gamma}_2| + \varepsilon_n/(b_n - a_n))^{1/4}} \geq e^{-c' n^{1/5}} \geq e^{-n\varepsilon_n^2} \qquad (13)$$

since $\varepsilon_n/(b_n - a_n) \to 0$ no faster than polynomially in n. Combining the above, we have that $\Pi_n(\|\bar{w}_n - W\|_{L^\infty([a_n,b_n])} \leq (3+\tilde{c})\varepsilon_n) \geq e^{-(2+c')n\varepsilon_n^2}$. □

Proof of Corollary 1. We shall use the notation employed in the proof of Theorem 3. By (9), it suffices to lower bound the prior probability of an L^∞-small ball about $\bar{w}_n$, where $\bar{f}_n = f_{\bar{w}_n}$ is the approximation (8). Note that since $\bar{N} \leq N$ (at least for n large enough), we can add additional breakpoints to the piecewise linear function $\bar{w}_n$ with weights $\bar{p}_i = 0$, $i = \bar{N}+1, ..., N$, without changing $\bar{w}_n$. Without loss of generality, pick any such additional breakpoints to be no smaller than $cn^{-6/5}$. Using similar computations to (10), for any $x \in [a_n, b_n]$,

$$\begin{aligned}
|\bar{w}_n(x) - w(x)| &= \left| \bar{\gamma}_1 \sum_{i=1}^N \frac{z_i \wedge (x-a_n)}{z_i} \bar{p}_i - \bar{\gamma}_2(x-a_n) - \gamma_1 \sum_{i=1}^N \frac{\theta_i \wedge (x-a_n)}{\theta_i} p_i - \gamma_2(x-a_n) \right| \\
&\leq |\bar{\gamma}_1 - \gamma_1| \sum_{i=1}^N \frac{\theta_i \wedge (x-a_n)}{\theta_i} p_i + \bar{\gamma}_1 \left| \sum_{i=1}^N \frac{\theta_i \wedge (x-a_n)}{\theta_i} p_i - \sum_{i=1}^N \frac{z_i \wedge (x-a_n)}{z_i} p_i \right| \\
&\quad + \bar{\gamma}_1 \sum_{i=1}^N \frac{z_i \wedge (x-a_n)}{z_i} |p_i - \bar{p}_i| + (b_n - a_n)|\bar{\gamma}_2 - \gamma_2| \\
&\leq |\bar{\gamma}_1 - \gamma_1| + \bar{\gamma}_1 \sum_{i=1}^N \frac{|\theta_i - z_i|}{\theta_i \wedge z_i} p_i + \bar{\gamma}_1 \sum_{i=1}^N |p_i - \bar{p}_i| + (b_n - a_n)|\bar{\gamma}_2 - \gamma_2|.
\end{aligned}$$

The first and fourth terms are bounded from above by ε_n with prior probability at least $e^{-n\varepsilon_n^2}$ by (12) and (13), respectively, for both priors. For the second term note, similarly to the proof of Theorem 3, that $z_1 \geq cn^{-6/5}$. Taking $\theta_i \in [z_i, z_i + cn^{-6/5}\varepsilon_n/\bar{\gamma}_1]$, the second term is bounded by $\bar{\gamma}_1 \sum_{i=1}^N p_i \varepsilon_n/\bar{\gamma}_1 = \varepsilon_n$. The probability of this set under the base measure is $H([z_i, z_i + cn^{-6/5}\varepsilon_n/\bar{\gamma}_1]) \gtrsim cn^{-6/5}\varepsilon_n/\big(\bar{\gamma}_1(b_n - a_n)^\eta\big)$ by the assumptions on $\bar{H}$. The joint probability that $\theta_i \in [z_i, z_i + cn^{-6/5}\varepsilon_n/\bar{\gamma}_1]$ for every $i = 1, ..., N$ is therefore bounded from below by a multiple of $(cn^{-6/5}\varepsilon_n/\bar{\gamma}_1)^N/(b_n - a_n)^{\eta N} \gtrsim e^{-c_1 N \log n} \geq e^{-c_2 n\varepsilon_n^2}$, for sufficiently large constants $c_1, c_2 > 0$.

It remains to show that the second term is bounded from above by ε_n with probability at least $e^{-cn\varepsilon_n^2}$ for some $c > 0$. In the case where $(p_1, ..., p_N)$ is endowed with a Dirichlet distribution, this statement follows from (11). In the case of the truncated stick-breaking prior, writing $(\bar{p}_1, \ldots, \bar{p}_N)$ in decreasing order, we note that there exist $0 \leq \bar{v}_1, ..., \bar{v}_{N-1} \leq 1$ such that $\bar{p}_i = \prod_{j=1}^{i-1}(1-\bar{v}_j)\bar{v}_i$, $i = 1, ..., N-1$, and $\bar{p}_N = \prod_{j=1}^{N-1}(1-\bar{v}_j)$. Define for $i = 1, 2, ..., N-1$ the intervals

$$I_i = \big[(\bar{v}_i - \varepsilon_n n^{-4/5} N^{-2}) \vee \varepsilon_n n^{-4/5} N^{-2}/2, (\bar{v}_i + \varepsilon_n n^{-4/5} N^{-2}) \wedge (1 - \varepsilon_n n^{-4/5} N^{-2}/2)\big].$$

For $v_i \in I_i \subset [0,1]$, $i = 1, .., N-1$, we have

$$\begin{aligned}
\Big|(1-v_1)...(1-v_i)v_{i+1} - (1-\bar{v}_1)...(1-\bar{v}_i)\bar{v}_{i+1}\Big| &\leq |v_1 - \bar{v}_1| + |v_2 - \bar{v}_2| + ... + |v_{i+1} - \bar{v}_{i+1}| \\
&\leq (i+1)\varepsilon_n n^{-4/5} N^{-2} \leq \varepsilon_n n^{-4/5} N^{-1}.
\end{aligned}$$

Hence for $p_i := (1-v_1)(1-v_2)...(1-v_{i-1})v_i$, $i = 1, ..., N-1$, $(p_1, .., p_N)$ is in the N-dimensional simplex and $\bar{\gamma}_1 \sum_{i=1}^N |\bar{p}_i - p_i| \leq \bar{\gamma}_1 N \varepsilon_n n^{-4/5} N^{-1} \lesssim \varepsilon_n$. Finally, we note that for $v_i \sim \text{Beta}(a, b)$, we have that $P(v_i \in I_i) \gtrsim (\varepsilon_n n^{-4/5} N^{-2})^{a \vee b}$ and we can therefore conclude

$$P\bigg(\bar{\gamma}_1 \sum_{i=1}^N |p_i - \bar{p}_i| \leq c\varepsilon_n\bigg) \geq \prod_{i=1}^{N-1} P(v_i \in I_i) \geq e^{N(a\vee b)\log(\varepsilon_n N^{-2} n^{-4/5})} \geq e^{-c_1 N \log n} \geq e^{-c_2 n \varepsilon_n^2},$$

for some large enough constants $c_1, c_2 > 0$, thereby completing the proof. □

5 Technical results

5.1 Proof of Proposition 1

In this section, we construct the piecewise log-linear approximation for an upper semi-continuous log-concave density given in Proposition 1. In particular, we require that the number of knots in the approximating function does not grow too quickly and that the knots are polynomially separated, thereby rendering the construction somewhat involved. The proof relies on firstly approximating any continuous concave function on a given compact interval using a piecewise linear function. One then splits $\text{supp}(f_0)$ into sets, depending on the size of both $\log f_0$ and $|(\log f_0)'|$, and obtains suitable piecewise linear approximations defined locally on each of these sets. Piecing together these local functions gives the desired global approximation.

We now construct a piecewise linear approximation of a continuous concave function w on a compact interval $[a, b]$. For any partition $a = x_0 < x_1 < \cdots < x_m = b$ of $[a, b]$, let $\widetilde{w}_m$ denote the piecewise linear approximation of w given by

$$\widetilde{w}_m(x) := \sum_{i=2}^m \bigg(\frac{x - x_{i-1}^*}{x_i^* - x_{i-1}^*} \frac{1}{x_i - x_{i-1}} \theta_i + \frac{x_i^* - x}{x_i^* - x_{i-1}^*} \frac{1}{x_{i-1} - x_{i-2}} \theta_{i-1} \bigg) \mathbb{1}_{(x_{i-1}^*, x_i^*]}(x), \tag{14}$$

where $\theta_i := \int_{x_{i-1}}^{x_i} w(s)ds$ and $x_i^* := \frac{x_i + x_{i-1}}{2}$. On $[a, x_1^*]$ and $(x_m^*, b]$, the function is defined by linearly extending the piecewise linear function defined above, that is

$$\begin{aligned} \widetilde{w}_m(a) &:= \frac{1}{x_2^* - x_1^*} \Big(\frac{x_2^* - a}{x_1 - a} \theta_1 - \frac{x_1^* - a}{x_2 - x_1} \theta_2 \Big), \\ \widetilde{w}_m(b) &:= \frac{1}{x_m^* - x_{m-1}^*} \Big(\frac{b - x_{m-1}^*}{b - x_{m-1}} \theta_m - \frac{b - x_m^*}{x_{m-1} - x_{m-2}} \theta_{m-1} \Big). \end{aligned} \tag{15}$$

The function $\widetilde{w}_m$ takes value $\widetilde{w}_m(x_i^*) = \frac{1}{x_i - x_{i-1}} \int_{x_{i-1}}^{x_i} w(s)ds$ at the midpoint $x_i^* = \frac{x_i + x_{i-1}}{2}$ of the interval $[x_{i-1}, x_i]$ and interpolates linearly in between.

Lemma 2. *Let $w : [a, b] \to \mathbb{R}$ be a continuous concave function, where $-\infty < a < b < \infty$. For any partition $a = x_0 < x_1 < \cdots < x_m = b$ of $[a, b]$, let $\widetilde{w}_m$ denote the piecewise linear approximation of w defined in (14) and (15). Then $\widetilde{w}_m$ is a concave function.*

Proof. By rescaling, we may without loss of generality assume that $[a,b] = [0,1]$. Note that $\widetilde{w}_m$ is concave if and only if

$$\widetilde{w}_m(x_i^*) \geq \frac{x_i^* - x_{i-1}^*}{x_{i+1}^* - x_{i-1}^*}\widetilde{w}_m(x_{i+1}^*) + \frac{x_{i+1}^* - x_i^*}{x_{i+1}^* - x_{i-1}^*}\widetilde{w}_m(x_{i-1}^*) \tag{16}$$

for $i = 2, \ldots, m-1$. Indeed, since $\widetilde{w}_m$ is piecewise linear, it is concave if and only if at every point where the derivative is discontinuous (i.e. a knot), the left derivative is greater than or equal to the right derivative. The above statement follows since the derivative of $\widetilde{w}_m$ is discontinuous (at most) at the points x_i^*, $i = 2, \ldots, m-1$, where the desired inequality is:

$$\frac{\widetilde{w}_m(x_{i+1}^*) - \widetilde{w}_m(x_i^*)}{x_{i+1}^* - x_i^*} \leq \frac{\widetilde{w}_m(x_i^*) - \widetilde{w}_m(x_{i-1}^*)}{x_i^* - x_{i-1}^*},$$

which is equivalent to (16).

To see that (16) holds if w is concave, we argue by contradiction and suppose that there exists i such that $\widetilde{w}_m(x_i^*) < \frac{x_i^* - x_{i-1}^*}{x_{i+1}^* - x_{i-1}^*}\widetilde{w}_m(x_{i+1}^*) + \frac{x_{i+1}^* - x_i^*}{x_{i+1}^* - x_{i-1}^*}\widetilde{w}_m(x_{i-1}^*)$. Consider the linear function l,

$$l(x) := \frac{x - x_{i-1}^*}{x_{i+1}^* - x_{i-1}^*}\widetilde{w}_m(x_{i+1}^*) + \frac{x_{i+1}^* - x}{x_{i+1}^* - x_{i-1}^*}\widetilde{w}_m(x_{i-1}^*).$$

In particular, we have that $\widetilde{w}_m(x_{i-1}^*) - l(x_{i-1}^*) = \widetilde{w}_m(x_{i+1}^*) - l(x_{i+1}^*) = 0$ and $\widetilde{w}_m(x_i^*) < l(x_i^*)$. We further denote $g := w - l$ and observe that

$$\tilde{g}_m(x_i^*) = \frac{1}{x_i - x_{i-1}}\int_{x_{i-1}}^{x_i} \big(w(s) - l(s)\big)ds = \widetilde{w}_m(x_i^*) - \int_{x_{i-1}}^{x_i} l(s)ds = \widetilde{w}_m(x_i^*) - l(x_i^*).$$

It follows that $\tilde{g}_m(x) = \widetilde{w}_m(x) - l(x)$ for all $x \in [0,1]$ and hence by the mean value theorem,

$$\tilde{g}_m(x_i^*) = \frac{1}{x_i - x_{i-1}}\int_{x_{i-1}}^{x_i} g(s)ds = g(\xi_i)$$

for some $\xi_i \in [x_{i-1}, x_i]$. One can similarly prove the existence of two points, $\xi_{i-1} \in [x_{i-2}, x_{i-1}]$ and $\xi_{i+1} \in [x_i, x_{i+1}]$, such that $\tilde{g}_m(x_{i-1}^*) = g(\xi_{i-1})$ and $\tilde{g}_m(x_{i+1}^*) = g(\xi_{i+1})$. Using the above results, we deduce the existence of three points $\xi_{i-1} < \xi_i < \xi_{i+1}$ such that $g(\xi_{i-1}) = 0 = g(\xi_{i+1})$ and $g(\xi_i) < 0$, which is a contradiction since g is concave by the concavity of w and l. □

Lemma 3. *Let $w : [a,b] \to \mathbb{R}$ be a continuous concave function with $w'_+(a) - w'_-(b) \leq M$ and where $-\infty < a < b < \infty$. Then there exists a partition $a = x_0 < x_1 < \cdots < x_m = b$ of $[a,b]$ with $\min_{i=1,\ldots,m}(x_i - x_{i-1}) \geq (b-a)(2m)^{-2}$ and such that*

$$\sup_{x \in [a,b]} |w(x) - \widetilde{w}_m(x)| \leq C\frac{M(b-a)}{m^2},$$

where $\widetilde{w}_m$ is the piecewise linear approximation of w defined in (14) *and* (15) *and $C > 0$ does not depend on a, b, m.*

Proof. By translation we may without loss of generality take $a = 0$. Recall that since w is a continuous concave function, it has left and right derivatives at every point $x \in [0, b]$. Define $\Delta w'(x) = w'_-(x) - w'_+(x)$. For every $r \geq 1$, let $\mathcal{P}_1 := \{x_{i,1} := \frac{ib}{r}, i = 0, \ldots, r\}$ be the uniform partition of $[0, b]$ and let $\tilde{x}_{i,2}$, $i = 1, \ldots, r_2$, be the points such that $\Delta w'(\tilde{x}_{i,2}) \geq M/r$, setting $r_2 = 0$ if no such point exists. By concavity of w,

$$M \geq w'_+(0) - w'_-(b) \geq \sum_{i=1}^{r_2} \Delta w'(\tilde{x}_{i,2}) \geq \frac{Mr_2}{r}, \tag{17}$$

so that $r_2 \leq r$.

Consider a new partition $\mathcal{P}_2 := \{x_{0,2} < \cdots < x_{r'_2,2}\}$ of $[0, b]$, consisting of the points $\{x_{i,1}\} \cup \{\tilde{x}_{i,2}\} \cup \{\tilde{x}_{i,2} - br^{-2}\} \cup \{\tilde{x}_{i,2} + br^{-2}\}$ written in increasing order. Note that $r'_2 \leq r + 3r_2$. Colour in red all the points of the form $\tilde{x}_{i,2}$ and $\tilde{x}_{i,2} - br^{-2}$, so that each red point is the left endpoint of an interval of length at most br^{-2}. This colouring will be used to keep track of points that have a close neighbour on the right.

We next refine the partition $\mathcal{P}_2$ by adding the point y between $x_{i,2}$ and $x_{i+1,2}$

$$y := \sup \left\{ x > x_{i,2} : w'_+(x_{i,2}) - w'_-(x) \leq \frac{2M}{r} \right\}$$

if

$$w'_+(x_{i,2}) - w'_-(x_{i+1,2}) > \frac{2M}{r}.$$

Denote by r_3 the total number of points y added in this manner to the sequence. We further add the points $y - br^{-2}$, $y + br^{-2}$ and colour in red all points of the form y and $y - br^{-2}$, similarly to the previous case. Repeating this procedure results in a new partition that separates intervals where the derivative decreases by at most $2M/r$. Denote by $\mathcal{P}_3 := \{0 = x_{0,3} < x_{1,3} < \cdots < x_{r'_3,3} = b\}$ this new partition. We now show that $r'_3 \leq 7r$.

Let y by any point added in the way just described. Suppose by contradiction that $w'_+(x_{i,2}) - w'_-(y) < M/r$. By definition, we know that for all $x > y$, $w'_+(x_{i,2}) - w'_-(x) > 2M/r$. Subtracting the two inequalities gives $w'_-(y) - w'_-(x) > M/r$. However, since the right derivative of a concave function is right continuous, taking the limit $x \to y^+$ (and restricting to the points x where w is differentiable) yields $\Delta w'(y) \geq M/r$. This is a contradiction however, because if this were the case, y would already belong to $\mathcal{P}_2$. Since $w'_+(x_{i,2}) - w'_-(y) \geq M/r$, using a similar argument to (17) gives $r_3 \leq r$ so that $r'_3 \leq 7r$.

Finally, if the function w is not differentiable at the point $x^*_{i,3} = \frac{x_{i,3} + x_{i-1,3}}{2}$, we split $[x_{i-1,3}, x_{i,3}]$ into two parts in such a way that w is differentiable at the midpoints of both new intervals and each interval has size at least $(x_{i,3} - x_{i-1,3})/3$. We add the points separating the new intervals to the previous partition, thereby obtaining $\mathcal{P}_4 := \{0 = x_{0,4} < x_{1,4} < \cdots < x_{\nu,4} = b\}$. The cardinality of $\mathcal{P}_4$, satisfies $\nu + 1 \leq 14r + 1$. We now create a new partition $\mathcal{P}$ with polynomially separated points using the following algorithm.

1. Set $\mathcal{P} = \mathcal{P}_4$, keeping track of all the points coloured red. Set $\tilde{x} = x_{0,4}$.

2. If $b - \tilde{x} \leq br^{-2}$, remove all points in $\mathcal{P}$ strictly between $\tilde{x}$ and b skip to Step 4.

3. Set $y = \inf\{t \in \mathcal{P} : t > \tilde{x} + br^{-2}\}$. Remove all elements of $\mathcal{P}$ between $\tilde{x}$ and y. If at least one element was removed, add to $\mathcal{P}$ the point $s = \tilde{x} + br^{-2} + \varepsilon$ for some $0 < \varepsilon < br^{-2} \wedge (y - \tilde{x} - br^{-2})$ such that w is differentiable at $(s+y)/2$ and $(s+\tilde{x})/2$. Colour $\tilde{x}$ red to mark that $s - \tilde{x} < 2br^{-2}$. Set $\tilde{x} := s$. If no point was removed from $\mathcal{P}$, set $\tilde{x} := y$. Go to Step 2.

4. If $\tilde{x} = b$ then stop. Otherwise set $y = \max\{t \in \mathcal{P} : t < \tilde{x}\}$ and remove $\tilde{x}$ from $\mathcal{P}$. If $b - y > 2br^{-2}$, add the point $s := b - br^{-2} - \varepsilon$ to $\mathcal{P}$ and colour it red, where $0 < \varepsilon < (b-y-2br^{-2}) \wedge br^{-2}$ is such that w is differentiable at $(y+s)/2$. If $b - y \le 2br^{-2}$, add the point $s := (y+b)/2$ and colour both y and s red.

Relabel the final partition $\mathcal{P} := \{0 = x_0 < x_1 < \cdots < x_m = b\}$ and note that

$$r \le m \le \nu \le 14r + 1.$$

By construction $\min_{i=0,\ldots,m-1}(x_{i+1} - x_i) \ge \frac{1}{2}br^{-2} \ge \frac{1}{2}bm^{-2}$ and

$$x_{i+1} - x_i \le \begin{cases} 2C_0^2 bm^{-2} & \text{if } x_i \text{ is coloured red (with } C_0 = 15), \\ C_0 bm^{-1} & \text{otherwise,} \end{cases}$$

since if x_i is coloured red, $x_{i+1} - x_i \le 2br^{-2} \le 2C_0^2 bm^{-2}$.

We now show that $\|w - \widetilde{w}_m\|_\infty = O(bm^{-2})$. If x_{i-1} is red, then by the mean value theorem, there exists $\xi_i \in J_i := [x_{i-1}, x_i]$ such that $\widetilde{w}_m(x_i^*) = w(\xi_i)$. Using the Lipschitz continuity of w and $\widetilde{w}_m$,

$$|w(x) - \widetilde{w}_m(x)| \le 2C_0^2 \frac{bM}{m^2}, \quad \forall x \in J_i := [x_{i-1}, x_i].$$

If x_{i-1} is not red, Taylor expanding w at the points $x_i^* := \frac{x_i + x_{i-1}}{2}$ (at which w is differentiable by the construction of $\mathcal{P}$) gives

$$w(x) = w(x_i^*) + w'(x_i^*)(x - x_i^*) + R_i(x), \quad x \in J_i. \tag{18}$$

Due to the construction of $\mathcal{P}$,

$$\begin{aligned} |R_i(x)| &= \left| w(x) - w(x_i^*) - w'(x_i^*)(x - x_i^*) \right| \\ &= \left| (w'(\xi_i) - w'(x_i^*))(x - x_i^*) \right|, \end{aligned}$$

where $w'(\xi_i)$ stands here for some value in the interval $[w'_+(\xi_i), w'_-(\xi_i)]$ for some point $\xi_i \in J_i$. We then deduce that $|R_i(x)| \le \frac{2M}{r}\frac{C_0 b}{2m} \le C_0^2 bM/m^2$.

Since $\widetilde{w}_m$ is piecewise linear, we can write

$$\widetilde{w}_m(x) = \widetilde{w}_m(x_i^*) + \widetilde{w}'_m(x_i^*)(x - x_i^*), \qquad x \in J_i,$$

where $\widetilde{w}'_m$ denotes the left or right derivative of $\widetilde{w}_m$ at x_i^*, depending on whether $x < x_i^*$ or $x > x_i^*$. We now show that $|\widetilde{w}'_m(x_i^*) - w'(x_i^*)| \le \max\{w'_+(x_{i-1}) - w'(x_i^*), w'(x_i^*) - w'_-(x_{i+1})\}$

for $i = 1, \dots, m-1$. Consider the case of right derivatives (the same argument also works for left derivatives). Using the definition of $\widetilde{w}_m$ and that $\theta_i = \int_{J_i} w(s)ds$,

$$\begin{aligned}\widetilde{w}'_{m,+}(x_i^*) &= \frac{\widetilde{w}_m(x_{i+1}^*) - \widetilde{w}_m(x_i^*)}{x_{i+1}^* - x_i^*} = \frac{1}{x_{i+1}^* - x_i^*}\left(\frac{\theta_{i+1}}{x_{i+1} - x_i} - \frac{\theta_i}{x_i - x_{i-1}}\right) \\ &= \frac{1}{(x_i - x_{i-1})(x_{i+1}^* - x_i^*)} \int_{x_{i-1}}^{x_i} \left[w\Big(\frac{x_{i+1} - x_i}{x_i - x_{i-1}} t + \frac{x_i^2 - x_{i+1}x_{i-1}}{x_i - x_{i-1}}\Big) - w(t)\right] dt\end{aligned}$$

and

$$\int_{x_{i-1}}^{x_i} \left[\frac{x_{i+1} - x_i}{x_i - x_{i-1}} t + \frac{x_i^2 - x_{i+1}x_{i-1}}{x_i - x_{i-1}} - t\right] dt = (x_{i+1}^* - x_i^*)(x_i - x_{i-1}).$$

By the continuity and concavity of w, $(v-u)w'_-(x_{i+1}) \le w(v) - w(u) \le (v-u)w'_+(x_{i-1})$ for any $x_{i-1} \le u \le v \le x_{i+1}$. Combining all of the above yields

$$w'_-(x_{i+1}) \le \widetilde{w}'_{m,+}(x_i^*) \le w'_+(x_{i-1}). \tag{19}$$

We remark that $\max\{w'_+(x_{i-1}) - w'(x_i^*), w'(x_i^*) - w'_-(x_{i+1})\} \le \frac{5M}{r}$. Indeed, since $x_i - x_{i-1} > C_0^2 b/m^2$, the point x_i is not equal to $x'_{j,4}$ for any j and hence both $\Delta w'(x_i) < M/r$ and $w'_+(x_i) - w'_-(x_{i+1}) \le 2M/r$ hold. Together with $w'_+(x_{i-1}) - w'(x_i^*) \le 2M/r$ and $w'(x_i^*) - w'_-(x_i) \le 2M/r$, this verifies the preceding statement. Then

$$w'(x_i^*) - w'_-(x_{i+1}) \le w'(x_i^*) - w'_-(x_i) + w'_-(x_i) - w'_+(x_i) + w'_+(x_i) - w'_-(x_{i+1}) \le \frac{5M}{r}.$$

We hence deduce that $|w'(x_i^*) - \widetilde{w}'_m(x_i^*)| \le 5MC_0 m^{-1}$. Finally, using (18) and the fact that $\int_{J_i}(x - x_i^*)dx = 0$,

$$\big|w(x_i^*) - \widetilde{w}_m(x_i^*)\big| = \frac{1}{x_i - x_{i-1}}\left|\int_{J_i} \big(w(x_i^*) - w(x)\big)dx\right| \le \sup_{x \in J_i} |R_i(x)| \le \frac{C_0^2 bM}{m^2}.$$

Collecting together all the pieces, we have that for any $x \in [x_{i-1}, x_i]$,

$$|w(x) - \widetilde{w}_m(x)| \le |w(x_i^*) - \widetilde{w}_m(x_i^*)| + |x - x_i^*||\widetilde{w}'_m(x_i^*) - w'(x_i^*)| + |R_i(x)| \le \frac{9}{2} C_0^2 Mbm^{-2}.$$

□

Lemma 4. *Any piecewise linear concave function $w : [a,b] \to \mathbb{R}$ with N knots $\{z_1, ..., z_N\}$ can be written in the form*

$$w(x) = \gamma_1 \sum_{i=1}^{N} \frac{z_i \wedge (x-a)}{z_i} p_i - \gamma_2(x-a) + \gamma_3,$$

with parameters $0 \le \gamma_1 \le (w'_+(a) - w'_-(b))(b-a)$, $|\gamma_2| \le |w'_-(b)|$, $\gamma_3 \in \mathbb{R}$, $\sum_{i=1}^N p_i = 1$ and $p_i \ge 0$ for $i = 1, \dots, N$.

Proof. The left derivative of w is a step function $g : (a, b] \mapsto \mathbb{R}$ with $g(a + \varepsilon) = w'_+(a)$ for sufficiently small $\varepsilon > 0$ and $g(b) = w'_-(b)$. By shifting this function vertically by $-w'_-(b)$, we arrive at a non-negative, bounded, monotone decreasing step function, which can therefore be written as a monotone decreasing probability density times a normalizing constant γ_1. It is easy to see that $\gamma_1 \le (b-a)(w'_+(a) - w'_-(b))$. The step function can therefore be represented as

$$g(x) = \gamma_1 \int_{x-a}^{b-a} \frac{1}{z} dP_N(z) + w'_-(b), \qquad x \in [a, b],$$

with P_N an atomic probability measure with N atoms on $[0, b-a]$ and γ_1 the normalizing constant. Integrating the step function g yields

$$\bar{w}_N(x) = \gamma_1 \int_0^{b-a} \frac{z \wedge (x-a)}{z} dP_N(z) + w'_-(b)(x-a) + C,$$

which is equal to w for an appropriately chosen constant $C > 0$. □

Proof of Proposition 1. Let $\phi(x) = \mathbb{1}_{[-1,1]}(x)$ and set $\phi_n(x) = \phi(x/s_n)$ for $s_n = \frac{4}{5\alpha} \log n$. Let $f_1 = f_{1,n} = f_0\phi_n / \int f_0\phi_n$ denote the log-concave density function supported on $[-s_n, s_n]$ and note that $|1 - \int f_0\phi_n| \le P_{f_0}[-s_n, s_n]^c \le 2e^\beta \alpha^{-1} n^{-4/5}$. Arguing as in (7), one has that $h^2(f_0, f_{1,n}) \le 12 e^\beta \alpha^{-1} n^{-4/5}$ for $n \ge (4e^\beta/\alpha)^{5/4}$.

We write $f_1 = e^{w_1}$ and construct the approximating function $\bar{f}_n$ according to the value of w_1 and its left and right derivatives $w'_{1,-}$ and $w'_{1,+}$. Let

$$\begin{aligned}
A_0^n &= \{x \in [a_n, b_n] : w_1(x) < -\tfrac{4}{5} \log n\},\\
A_1^n &= \{x \in [a_n, b_n] : w_1(x) \ge -\tfrac{4}{5} \log n, |w'_{1,\pm}(x)| > n^{4/5}\},\\
A_{2,j}^n &= \{x \in [a_n, b_n] : w_1(x) \ge -\tfrac{4}{5} \log n, 2^{-j-1} n^{4/5} < |w'_{1,\pm}(x)| \le 2^{-j} n^{4/5}\}, \quad j = 0, \dots, j_n,\\
A_3^n &= \{x \in [a_n, b_n] : w_1(x) \ge -\tfrac{4}{5} \log n, |w'_{1,\pm}(x)| \le D\},
\end{aligned}$$

where $D > 0$ is some fixed constant, $|w'_{1,\pm}(x)| = \max(|w'_{1,+}(x)|, |w'_{1,-}(x)|)$ and $j_n = \lceil \log_2(n^{4/5}/D) - 1 \rceil$. In fact the set where the left and right derivatives of the concave function w_1 do not agree has measure zero. Note that the above sets are all disjoint except A_{2,j_n}^n and A_3^n: since j_n is the smallest integer such that $2^{-j_n - 1} n^{4/5} \le D$, these last two sets may overlap. In particular, we can express $[a_n, b_n]$ as the almost disjoint union of the above sets. Write $B_n = (\cup_{j=0}^{j_n} A_{2,j}^n) \cup A_3^n \subset [a_n, b_n]$ for convenience and note that by the concavity of w_1, this is an interval. Since $\|f_1\|_\infty \le e^{2\beta}$ for $n \ge (4e^\beta/\alpha)^{5/4}$, the set A_1^n consists of two intervals, each of width $O(n^{-4/5} \log n)$. Using again the boundedness of f_1, the definition of A_0^n and that $|\mathrm{supp}(f_1)| \lesssim \log n$,

$$\int_{B_n} f_1 dx = 1 - O(n^{-4/5} \log n), \tag{20}$$

so that in particular, $B_n \ne \emptyset$ for $n \ge n_0(\alpha, \beta)$ large enough.

For an interval $I = [a, b]$, let $\mathcal{P}_{n,I}$ denote the corresponding partition given by Lemma 3 with partition size $\lceil n^{1/5} + 1 \rceil$ (rather than $m + 1$ as in the statement of the lemma). The piecewise linear approximation (14) of the function w_1 based on $\mathcal{P}_{n,I}$ will therefore satisfy

the conclusions of Lemma 3 on I. Note that $A^n_{2,j}$ consists of at most two disjoint intervals, which by the boundedness of f_1 are each of length $O(2^{j+1}n^{-4/5}\log n)$. In a slight abuse of notation, denote by $\mathcal{P}_{n,A^n_{2,j}}$ the union of the partitions from the two disjoint intervals, thereby placing $\lceil n^{1/5}+1\rceil$ points in each. Similarly, A^n_3 consists of a single interval of length $O(\log n)$. Consider now the overall partition $\mathcal{P}_n = (\cup_{j=0}^{j_n-1}\mathcal{P}_{n,A^n_{2,j}}) \cup \mathcal{P}_{n,A^n_{2,j_n}\cap(A^n_3)^c} \cup \mathcal{P}_{n,A^n_3}$ of B_n, which has $O(j_n n^{1/5}) = O(n^{1/5}\log n)$ points. Note that we are putting points more densely in $A^n_{2,j}$ for small j, since these intervals are smaller. The associated piecewise linear function $\widetilde{w}_n$ defined in (14) and based on $\mathcal{P}_n$ is concave by Lemma 2 and by construction corresponds to the partition given in Lemma 3 for each of the sets comprising B_n. It therefore satisfies:

- $\|w_1 - \widetilde{w}_n\|_{L^\infty(A^n_{2,j})} \le Cn^{-2/5}\log n$ for some $C>0$ independent of j,
- the partition points in the set $A^n_{2,j}$ are distance at least $c2^j n^{-6/5}\log n \ge cn^{-6/5}\log n$ apart, for some $c>0$ independent of j.

For A^n_3, the function satisfies the same L^∞-bound with the partition points being $n^{-2/5}\log n$-separated. Finally note that since these intervals meet only at their boundaries, and the boundary points of the intervals are contained in the partition presented in Lemma 3, the interval boundaries will be contained in $\mathcal{P}_n$. Consequently, the separation property continues to hold even across the different subpartitions. In conclusion, we have shown that $\widetilde{w}_n$ is concave and piecewise linear with $O(n^{1/5}\log n)$ knots, which are $c'n^{-6/5}$-separated, and satisfies

$$\sup_{x\in B_n} |\widetilde{w}_n(x) - w_1(x)| = O(n^{-2/5}\log n). \tag{21}$$

We now extend the approximating function to $[a_n,b_n]\backslash B_n$. Write $\mathcal{P}_n = (x_i)_{i=0}^M$, where $\min(B_n) = x_0 < x_1 < ... < x_M = \max(B_n)$ and $M = O(n^{1/5}\log n)$. Define $\bar{w}_n : [a_n,b_n]\to\mathbb{R}$ as

$$\bar{w}_n(x) = \begin{cases} \widetilde{w}_n(x_0) + \big(w'_{1,-}(x_0)\wedge n^{4/5}\vee(-n^{4/5})\big)(x-x_0) & x\in[a_n,x_0], \\ \widetilde{w}_n(x) & x\in B_n, \\ \widetilde{w}_n(x_M) + \big(w'_{1,+}(x_M)\wedge n^{4/5}\vee(-n^{4/5})\big)(x-x_M) & x\in[x_M,b_n]. \end{cases} \tag{22}$$

This is simply the function $\widetilde{w}_n$ extended linearly from the boundary points of B_n with slope $w'_{1,-}(x_0)\wedge n^{4/5}\vee(-n^{4/5})$ and $w'_{1,+}(x_M)\wedge n^{4/5}\vee(-n^{4/5})$ on $[a_n,x_0]$ and $[x_M,b_n]$ respectively. We now verify that $\bar{w}_n$ is concave, for which it is enough to show that $\bar{w}'_{n,+}(x_0)\le\bar{w}'_{n,-}(x_0)$ and $\bar{w}'_{n,+}(x_M)\le\bar{w}'_{n,-}(x_M)$. For the first inequality, using (19), the concavity of w_1 and the boundary construction of $\widetilde{w}_n$ given by (14), $\bar{w}'_{n,+}(x_0) = \widetilde{w}'_{n,+}(x_0) = \widetilde{w}'_{n,+}(x_1^*) \le w'_{1,+}(x_0) \le w'_{1,-}(x_0)$. Since $x_0\in B_n$, it also holds that $|\bar{w}'_{n,+}(x_0)|\le n^{4/5}$. The second inequality can be proved analogously.

Since $\log(1+z) = O(z)$ as $|z|\to 0$, it follows that $\log\int f_0\phi_n = O(n^{-4/5})$. Using this and (21),

$$|\bar{w}_n(x_0) - \log f_0(x_0)| \le \left|\log\int f_0\phi_n\right| + O(n^{-2/5}\log n) = O(n^{-2/5}\log n).$$

By concavity, the slope of the linear extension on $[a_n,x_0]$ satisfies $w'_{1,-}(x_0) = (\log f_0)'_-(x_0) \le (\log f_0)'_+(x)$ for all $x<x_0$ such that $f_0(x)>0$. Combining the above yields $\bar{w}_n(x)\ge$

$\log f_0(x) - O(n^{-2/5}\log n)$ for all $x \in [a_n, x_0]$. The same computation also gives the result for $x \in [x_M, b_n]$, so that for some $C > 0$,

$$\sup_{x\in[a_n,b_n]\setminus B_n} (\log f_0(x) - \bar{w}_n(x)) \le Cn^{-2/5}\log n. \tag{23}$$

Define the log-concave density

$$\bar{f}_n(x) = \begin{cases} e^{\bar{w}_n(x)} / \int_{a_n}^{b_n} e^{\bar{w}_n} & x \in [a_n, b_n], \\ 0 & x \notin [a_n, b_n]. \end{cases}$$

This function is piecewise log-linear, has $O(n^{1/5}\log n)$ knots and satisfies (ii) and (iii) by construction. We have

$$h^2(f_1, \bar{f}_n) \le 2\int_{B_n^c} f_1 + 2\int_{B_n^c} \bar{f}_n + \int_{B_n} (f_1^{1/2} - \bar{f}_n^{1/2})^2. \tag{24}$$

The first integral is $O(n^{-4/5}\log n)$ by (20). Using (21), $\int_{B_n} e^{\bar{w}_n} = e^{o(1)}\int_{B_n} f_1$. Notice that there are three cases for the end point x_0 of B_n: $x_0 \in \overline{A_0^n}$, $x_0 \in \overline{A_1^n}$ or $x_0 = a_n$. If $x_0 \in \overline{A_0^n}$, then since w_1 is increasing on $[a_n, x_0]$, $\int_{a_n}^{x_0} e^{\bar{w}_n} \le \int_{a_n}^{x_0} e^{-(4/5)\log n + n^{-2/5}\log n} \le (x_0 - a_n)e^{o(1)}n^{-4/5}$. If $x_0 \in \overline{A_1^n}$, then $\int_{a_n}^{x_0} e^{\bar{w}_n} \le \int_{a_n}^{x_0} e^{w_1(x_0) + n^{-2/5}\log n + n^{4/5}(x-x_0)}dx \le e^{2\beta + o(1)}n^{-4/5}$ for $n \ge (4e^{\beta}/\alpha)^{5/4}$. If $x_0 = a_n$, the bound is trivial. Using the same bounds for $[x_M, b_n]$ gives $\int_{B_n^c} e^{\bar{w}_n} = O((b_n - a_n)n^{-4/5})$, so that $\int_{a_n}^{b_n} e^{\bar{w}_n} = 1 + o(1)$. This implies that the second integral in (24) is $O((b_n - a_n)n^{-4/5})$.

Using (20), (21), Lemma 3.1 of [36] and the above, the third term of (24) is bounded by a multiple of

$$\begin{aligned} &\int_{B_n} e^{w_1}\left(1 - \frac{1}{\sqrt{\int_{B_n} e^{w_1}}}\right)^2 + \int_{B_n}\left(\frac{e^{w_1/2}}{\sqrt{\int_{B_n} e^{w_1}}} - \frac{e^{\bar{w}_n/2}}{\sqrt{\int_{B_n} e^{\bar{w}_n}}}\right)^2 \\ &\qquad + \int_{B_n} e^{\bar{w}_n}\left(\frac{1}{\sqrt{\int_{B_n} e^{\bar{w}_n}}} - 1\right)^2 \\ &\lesssim \left(\int_{B_n} f_1 - 1\right)^2 + \sup_{x\in B_n} |\bar{w}_n(x) - w_1(x)|^2 + \left(\int_{B_n} e^{\bar{w}_n} - 1\right)^2 \\ &= O((\log n)^2 n^{-4/5} + (b_n - a_n)^2 n^{-8/5}), \end{aligned}$$

which establishes (i).

Consider (iv). Note that this is trivial if $f_0(x) = 0$, so assume $f_0(x) \ne 0$. If $x \in B_n$, then by (21),

$$f_0(x)/\bar{f}_n(x) = e^{w_1(x) - \bar{w}_n(x)} \int f_0\phi_n \int e^{\bar{w}_n} = e^{O(n^{-2/5}\log n)}(1 + o(1)) = 1 + o(1).$$

If $x \in [a_n, b_n]\setminus B_n$, then the result follows from (23).

Consider lastly (v). Since $\bar{w}_n$ defined in (22) is piecewise linear with $|w'_+(a_n)| \vee |w'_-(b_n)| \leq n^{4/5}$, in view of Lemma 4 it takes the form

$$\bar{w}_n(x) = \gamma_1 \sum_{i=1}^{M} \frac{z_i \wedge (x - a_n)}{z_i} - \gamma_2(x - a_n) + \gamma_3, \quad x \in [a_n, b_n],$$

with $M = O(n^{1/5} \log n)$, $\gamma_1 \leq |w'_+(a_n) - w'_-(b_n)|(b_n - a_n) \leq 2n^{4/5}(b_n - a_n)$ and $|\gamma_2| \leq |w'_-(b_n)| \leq n^{4/5}$. This completes the proof. □

5.2 Proof of Lemma 1

Proof of Lemma 1. It is shown in the proof of Theorem 5 of Kim and Samworth [21], p. 2772, that for $\eta \in (0,1)$ to be defined below and $t \geq t_0(\eta)$,

$$\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n \left(\left\{ h(\hat{g}_n, g_0) \geq t n^{-2/5} \right\} \cap \left\{ \hat{g}_n \in \widetilde{\mathcal{F}}^{1,\eta} \right\} \right) \leq 2^{15/2} \exp\left(-\frac{t^2 n^{1/5}}{2^{28}} \right),$$

so that it remains only to control $\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(\hat{g}_n \notin \widetilde{\mathcal{F}}^{1,\eta})$. Lemma 6 of Kim and Samworth [21] shows that this quantity is $O(n^{-1})$ as $n \to \infty$; we essentially follow their proof, suitably sharpening the probability bounds in the case $d = 1$. Then

$$\begin{aligned} \sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(\hat{g}_n \notin \widetilde{\mathcal{F}}^{1,\eta}) \leq & \sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(|\mu_{\hat{g}_n}| > 1) + \sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(\sigma^2_{\hat{g}_n} > 1 + \eta) \\ & + \sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(\sigma^2_{\hat{g}_n} < 1 - \eta). \end{aligned} \tag{25}$$

By Lemma 13 of [12], there exist universal constants $\alpha_0, \beta_0 > 0$ such that for all $x \in \mathbb{R}$,

$$\sup_{g \in \overline{\mathcal{F}}^{0,1}} g(x) \leq e^{\beta_0 - \alpha_0 |x|}.$$

It is shown in [21], p. 2774, that the first term in (25) is bounded by $2\alpha_0^{-1} e^{\beta_0 - \alpha_0 \sqrt{n}}$. For the second term, by Remark 2.3 of Dümbgen et al. [11], one has $\sigma^2_{\hat{g}_n} \leq \widetilde{\sigma}^2_n := n^{-1} \sum_{i=1}^n (Z_i - \bar{Z}_n)^2 \leq n^{-1} \sum_{i=1}^n Z_i^2$, where $\widetilde{\sigma}^2_n$ denotes the sample covariance and $\bar{Z}_n = n^{-1} \sum_{i=1}^n Z_i$ the sample mean. Letting $C_0 = C_0(\alpha_0, \beta_0, 2)$ denote the constant in Lemma 5, we can apply that lemma to bound the second term by

$$\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n \left(\frac{1}{n} \sum_{i=1}^n Z_i^2 > 1 + \eta \right) \leq \exp(-\sqrt{\eta / C_0} n^{1/4})$$

for $n \geq \max(16 C_0^2 / \eta, e^2)$.

Consider now the third term in (25). Let $\overline{\mathcal{P}}_\varepsilon^{1/10, 1/2}$ denote the class of probability distributions on $\mathbb{R}$ such that $\mu_P = \int x dP(x)$ and $\sigma_P^2 = \int (x - \mu_P)^2 dP(x)$ satisfy $|\mu_P| \leq 1/10$ and $1/2 \leq \sigma_P^2 \leq 3/2$ and

$$\int |x|^{2+\varepsilon} dP(x) \leq 4 e^{\beta_0} \alpha_0^{-3-\varepsilon} \Gamma(3 + \varepsilon) =: \tau_\varepsilon,$$

where $\Gamma(t) = \int_0^\infty x^{t-1}e^{-x}dx$. This is exactly the same as the class $\mathcal{P}^{1/10,1/2}$ considered in Lemma 6 of [21], except that we have replaced the 4^{th}-moment condition with a $(2+\varepsilon)$-moment condition (the former condition is satisfied by the MLE with probability strictly larger than the desired probability). Following the rest of the proof of Lemma 6 of [21] (noting that $\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} \int |x|^{2+\varepsilon} g_0(x)dx \leq \tau_\varepsilon/2$ and that uniform integrability of $\{Y_{n_k}^2 : k \in \mathbb{N}\}$ in that proof follows from the $(2+\varepsilon)$-moment condition), one can similarly conclude that for some $\eta = \eta\big(\alpha_0, \beta_0, \varepsilon\big) \in (0,1)$,

$$\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(\sigma_{\hat{g}_n}^2 < 1-\eta) \leq \sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n\left(\mathbb{P}_n \notin \overline{\mathcal{P}}_\varepsilon^{1/10,1/2}\right),$$

where $\mathbb{P}_n = n^{-1}\sum_{i=1}^n \delta_{Z_i}$ is the empirical measure and δ_x is the Dirac measure at x. This last probability can be bounded by

$$\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(|\bar{Z}_n| > 1/10) + \sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(|\widetilde{\sigma}_n^2 - 1| > 1/2)$$
$$+ \sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n\left(\int |x|^{2+\varepsilon}(d\mathbb{P}_n(x) - g_0(x)dx) > \tau_\varepsilon/2\right).$$

The first two terms can be bounded using similar arguments to those used previously. For the last term, using Lemma 5 with $C_0 = C_0(\alpha_0, \beta_0, 2+\varepsilon)$ the constant in that lemma,

$$\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n\left(\frac{1}{n}\sum_{i=1}^n \left(|Z_i|^{2+\varepsilon} - \mathbb{E}_{g_0}|Z_i|^{2+\varepsilon}\right) > \frac{\tau_\varepsilon}{2}\right) \leq \exp\left(-\left(\frac{\tau_\varepsilon}{2C_0}\right)^{1/(2+\varepsilon)} n^{\frac{1}{4+2\varepsilon}}\right)$$

for $n \geq \max(4(2+\varepsilon)^{4+2\varepsilon}C_0^2/\tau_\varepsilon^2, e^{2+\varepsilon})$. In conclusion, we have shown that for $0 < \varepsilon < 1/2$ and $n \geq n_0(\varepsilon)$,

$$\sup_{g_0 \in \overline{\mathcal{F}}^{0,1}} P_{g_0}^n(\hat{g}_n \notin \widetilde{\mathcal{F}}^{1,\eta}) \leq C(\varepsilon)\exp(-c(\varepsilon)n^{\frac{1}{4+2\varepsilon}}).$$

This completes the proof. □

Lemma 5. *Let $Z_1, ..., Z_n$ be i.i.d. random variables from a density $f \in \mathcal{F}_{\alpha,\beta}$, where $\alpha > 0$ and $\beta \in \mathbb{R}$. Then for any $r \geq 1$, $t \geq r^r$ and $n \geq e^r$,*

$$P_f^n\left(\frac{1}{\sqrt{n}}\left|\sum_{i=1}^n (|Z_i|^r - \mathbb{E}|Z_i|^r)\right| \geq C_0(\alpha, \beta, r)t\right) \leq \exp(-t^{1/r}).$$

Proof. For notational convenience, write $\lambda = 1/r \in (0,1]$. Let $x_\lambda = (1/\lambda)^{1/\lambda}$ and define

$$\Psi_\lambda(x) = \begin{cases} e^{x^\lambda} & x \geq x_\lambda, \\ \tau_\lambda x & x < x_\lambda, \end{cases}$$

where $\tau_\lambda = \Psi_\lambda(x_\lambda)/x_\lambda = (\lambda e)^{1/\lambda}$. This defines a Young function, that is a convex, increasing function $\Psi_\lambda : \mathbb{R}^+ \to \mathbb{R}^+$ with $\Psi_\lambda(0) = 0$. Denote the corresponding Orlicz norm $\|X\|_{\Psi_\lambda} :=$

$\inf\{a > 0 : \mathbb{E}[\Psi_\lambda(|X|/a)] \le 1\}$. Note that the density function g_λ of $|Z_1|^{1/\lambda}$ satisfies $g_\lambda(x) = \lambda x^{-(1-\lambda)}(f(x^\lambda) + f(-x^\lambda))\mathbb{1}_{\{x\ge 0\}} \le 2\lambda x^{-(1-\lambda)} e^{\beta - \alpha x^\lambda}\mathbb{1}_{\{x\ge 0\}}$. Then for fixed $a > \alpha^{-1/\lambda}$,

$$\begin{aligned}\mathbb{E}\Psi_\lambda(|Z_1|^{1/\lambda}/a) &\le \frac{2\lambda\tau_\lambda e^\beta}{a}\int_0^{x_\lambda} u^\lambda e^{-\alpha u^\lambda} du + 2\lambda e^\beta \int_{x_\lambda}^\infty u^{-(1-\lambda)} e^{-(\alpha - a^{-\lambda})u^\lambda} du \\ &= K_0(a,\alpha,\beta,\lambda) < \infty.\end{aligned}$$

If $K_0 = K_0(a,\alpha,\beta,\lambda) > 1$, it follows by convexity that $\||Z_1|^{1/\lambda}\|_{\Psi_\lambda} \le aK_0$.

By Theorem 6.21 of Ledoux and Talagrand [23], there exists a constant K_λ such that

$$\left\|\sum_{i=1}^n (|Z_i|^{\frac{1}{\lambda}} - \mathbb{E}|Z_i|^{\frac{1}{\lambda}})\right\|_{\Psi_\lambda} \le K_\lambda \left(\left\|\sum_{i=1}^n (|Z_i|^{\frac{1}{\lambda}} - \mathbb{E}|Z_i|^{\frac{1}{\lambda}})\right\|_1 + \left\|\max_{1\le i\le n} \left||Z_i|^{\frac{1}{\lambda}} - \mathbb{E}|Z_i|^{\frac{1}{\lambda}}\right|\right\|_{\Psi_\lambda}\right).$$

The first-term on the right-hand side can be bounded by the $\|\cdot\|_2$-norm of the same quantity, which equals the square root of $n\mathbb{E}(|Z_i|^{\frac{1}{\lambda}} - \mathbb{E}|Z_i|^{\frac{1}{\lambda}})^2 \le n\mathbb{E}|Z_i|^{\frac{2}{\lambda}}$. For any $\delta \in (0,1]$,

$$\mathbb{E}|Z_1|^{\frac{1}{\delta}} = \int_0^\infty x g_\delta(x) dx \le 2\delta e^\beta \int_0^\infty x^\delta e^{-\alpha x^\delta} dx = \frac{2e^\beta}{\alpha^{1+\frac{1}{\delta}}}\int_0^\infty y^{\frac{1}{\delta}} e^{-y} dy = \frac{2e^\beta}{\alpha^{1+\frac{1}{\delta}}}\Gamma(1+1/\delta).$$

Since this is finite for any $\delta \in (0,1]$, we can bound the $\|\cdot\|_1$-norm above by $C(\alpha,\beta,\lambda)\sqrt{n}$.

Note that for any random variable X, $\|X - \mathbb{E}X\|_{\Psi_\lambda} \le 2\|X\|_{\Psi_\lambda}$. Indeed, setting $a = \|X\|_{\Psi_\lambda}$, since Ψ_λ is convex and increasing,

$$\mathbb{E}\Psi_\lambda\left(\frac{|X - \mathbb{E}X|}{2a}\right) \le \frac{1}{2}\mathbb{E}\Psi_\lambda\left(\frac{|X|}{a}\right) + \frac{1}{2}\mathbb{E}\Psi_\lambda\left(\frac{\mathbb{E}|X|}{a}\right) \le \mathbb{E}\Psi_\lambda\left(\frac{|X|}{a}\right) \le 1.$$

Using this and Lemma 2.2.2 of van der Vaart and Wellner [37],

$$\left\|\max_{1\le i\le n}\left||Z_i|^{\frac{1}{\lambda}} - \mathbb{E}|Z_i|^{\frac{1}{\lambda}}\right|\right\|_{\Psi_\lambda} \le K(\Psi_\lambda)\Psi_\lambda^{-1}(n)\max_{1\le i\le n}\left\||Z_i|^{\frac{1}{\lambda}}\right\|_{\Psi_\lambda} \le K(\alpha,\beta,\lambda)(\log n)^{\frac{1}{\lambda}},$$

so that we have shown

$$\left\|\sum_{i=1}^n (|Z_i|^{\frac{1}{\lambda}} - \mathbb{E}|Z_i|^{\frac{1}{\lambda}})\right\|_{\Psi_\lambda} \le C(\alpha,\beta,\lambda)\sqrt{n}.$$

By Markov's inequality, for any random variable X, $\mathbb{P}(|X| \ge x\|X\|_{\Psi_\lambda}) = \mathbb{P}(\Psi_\lambda(|X|/\|X\|_{\Psi_\lambda}) \ge \Psi_\lambda(x)) \le 1/\Psi_\lambda(x)$. Applying this to the above sum completes the proof. □

Acknowledgements: The authors would like to thank Richard Samworth and Arlene Kim for helpful discussions.

References

[1] Bagnoli, M., and Bergstrom, T. Log-concave probability and its applications. *Econom. Theory 26*, 2 (2005), 445–469.

[2] Balabdaoui, F., and Doss, C. R. Inference for a two-component mixture of symmetric distributions under log-concavity. *Bernoulli*, To appear.
[3] Brooks, S. P. MCMC convergence diagnosis via multivariate bounds on log-concave densities. *Ann. Statist. 26*, 1 (1998), 398–433.
[4] Caplin, A., and Nalebuff, B. Aggregation and imperfect competition: on the existence of equilibrium. *Econometrica 59*, 1 (1991), 25–59.
[5] Caplin, A., and Nalebuff, B. Aggregation and social choice: a mean voter theorem. *Econometrica 59*, 1 (1991), 1–23.
[6] Cule, M., and Samworth, R. Theoretical properties of the log-concave maximum likelihood estimator of a multidimensional density. *Electron. J. Stat. 4* (2010), 254–270.
[7] Cule, M., Samworth, R., and Stewart, M. Maximum likelihood estimation of a multi-dimensional log-concave density. *J. R. Stat. Soc. Ser. B Stat. Methodol. 72*, 5 (2010), 545–607.
[8] Doss, C. R., and Wellner, J. A. Inference for the mode of a log-concave density. *ArXiv e-prints* (Nov. 2016).
[9] Doss, C. R., and Wellner, J. A. Mode-constrained estimation of a log-concave density. *ArXiv e-prints* (Nov. 2016).
[10] Dümbgen, L., and Rufibach, K. Maximum likelihood estimation of a log-concave density and its distribution function: basic properties and uniform consistency. *Bernoulli 15*, 1 (2009), 40–68.
[11] Dümbgen, L., Samworth, R., and Schuhmacher, D. Approximation by log-concave distributions, with applications to regression. *Ann. Statist. 39*, 2 (2011), 702–730.
[12] Fresen, D. A multivariate Gnedenko law of large numbers. *Ann. Probab. 41*, 5 (2013), 3051–3080.
[13] Ghosal, S., Ghosh, J. K., and van der Vaart, A. W. Convergence rates of posterior distributions. *Ann. Statist. 28*, 2 (2000), 500–531.
[14] Ghosal, S., and van der Vaart, A. Posterior convergence rates of Dirichlet mixtures at smooth densities. *Ann. Statist. 35*, 2 (2007), 697–723.
[15] Ghosal, S., and van der Vaart, A. W. *Fundamentals of Nonparametric Bayesian Inference*.
[16] Giné, E., and Nickl, R. Rates of contraction for posterior distributions in L^r-metrics, $1 \leq r \leq \infty$. *Ann. Statist. 39*, 6 (2011), 2883–2911.
[17] Groeneboom, P., and Jongbloed, G. *Nonparametric estimation under shape constraints*, vol. 38 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2014. Estimators, algorithms and asymptotics.
[18] Hannah, L. A., and Dunson, D. B. Bayesian nonparametric multivariate convex regression. *ArXiv e-prints* (Sept. 2011).
[19] Khazaei, S., and Rousseau, J. Bayesian Nonparametric Inference of decreasing densities. In *42èmes Journées de Statistique* (Marseille, France, 2010).
[20] Kim, A. K. H., Guntuboyina, A., and Samworth, R. J. Adaptation in log-concave density estimation. *ArXiv e-prints* (2016).
[21] Kim, A. K. H., and Samworth, R. J. Global rates of convergence in log-concave density estimation. *Ann. Statist. 44*, 6 (2016), 2756–2779.
[22] Le Cam, L. *Asymptotic methods in statistical decision theory*. Springer Series in Statis-

tics. Springer-Verlag, New York, 1986.

[23] Ledoux, M., and Talagrand, M. *Probability in Banach spaces*. Classics in Mathematics. Springer-Verlag, Berlin, 2011. Isoperimetry and processes, Reprint of the 1991 edition.

[24] Lenk, P. J. The logistic normal distribution for Bayesian, nonparametric, predictive densities. *J. Amer. Statist. Assoc. 83*, 402 (1988), 509–516.

[25] Leonard, T. Density estimation, stochastic processes and prior information. *J. Roy. Statist. Soc. Ser. B 40*, 2 (1978), 113–146. With discussion.

[26] Mengersen, K. L., and Tweedie, R. L. Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Statist. 24*, 1 (1996), 101–121.

[27] Müller, S., and Rufibach, K. Smooth tail-index estimation. *J. Stat. Comput. Simul. 79*, 9-10 (2009), 1155–1167.

[28] Ray, K. Bayesian inverse problems with non-conjugate priors. *Electron. J. Stat. 7* (2013), 2516–2549.

[29] Salomond, J.-B. Concentration rate and consistency of the posterior distribution for selected priors under monotonicity constraints. *Electron. J. Stat. 8*, 1 (2014), 1380–1404.

[30] Samworth, R. J., and Yuan, M. Independent component analysis via nonparametric maximum likelihood estimation. *Ann. Statist. 40*, 6 (2012), 2973–3002.

[31] Seregin, A., and Wellner, J. A. Nonparametric estimation of multivariate convex-transformed densities. *Ann. Statist. 38*, 6 (2010), 3751–3781. With supplementary material available online.

[32] Shively, T. S., Sager, T. W., and Walker, S. G. A Bayesian approach to non-parametric monotone function estimation. *J. R. Stat. Soc. Ser. B Stat. Methodol. 71*, 1 (2009), 159–175.

[33] Shively, T. S., Walker, S. G., and Damien, P. Nonparametric function estimation subject to monotonicity, convexity and other shape constraints. *J. Econometrics 161*, 2 (2011), 166–181.

[34] Szabó, B., van der Vaart, A. W., and van Zanten, J. H. Frequentist coverage of adaptive nonparametric Bayesian credible sets. *Ann. Statist. 43*, 4 (2015), 1391–1428.

[35] van de Geer, S. A. *Applications of empirical process theory*, vol. 6 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2000.

[36] van der Vaart, A. W., and van Zanten, J. H. Rates of contraction of posterior distributions based on gaussian process priors. *Ann. Statist. 36*, 3 (2008), 1435–1463.

[37] van der Vaart, A. W., and Wellner, J. A. *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York, 1996. With applications to statistics.

[38] Walther, G. Detecting the presence of mixing with multiscale maximum likelihood. *J. Amer. Statist. Assoc. 97*, 458 (2002), 508–513.

[39] Walther, G. Inference and modeling with log-concave distributions. *Statist. Sci. 24*, 3 (2009), 319–327.

[40] Williamson, R. E. Multiply monotone functions and their Laplace transforms. *Duke Math. J. 23* (1956), 189–207.